

Distal and non-distal NIP theories

Pierre Simon
École Normale Supérieure, Paris

January 14, 2013

Abstract

We study one way in which stable phenomena can exist in an *NIP* theory. We start by defining a notion of ‘pure instability’ that we call ‘distality’ in which no such phenomenon occurs. O-minimal theories and the p-adics for example are distal. Next, we try to understand what happens when distality fails. Given a type p over a sufficiently saturated model, we extract, in some sense, the stable part of p and define a notion of stable-independence which is implied by non-forking and has bounded weight.

Contents

1	Introduction	2
1.0.1	A word about measures	3
1.0.2	Organization of the paper and main results	4
1.1	Preliminaries	6
1.1.1	Measures	7
1.1.2	Indiscernible sequences and cuts	8
1.2	Shrinking	10
1.3	Expanding	12
1.4	Sliding	12
1.5	Weight and dp-minimality	14

2	Distal theories	15
2.1	Indiscernible sequences	15
2.2	Invariant types	20
2.3	Generically stable measures	21
2.4	Reduction to dimension 1	25
3	Domination in non-distal theories	28
3.0.1	External characterization and base change	30
3.1	Domination for types	32
3.1.1	The moving-away lemma	33
3.1.2	S-independence	35
3.2	The finite-co-finite theorem and application	37
4	Sharp theories	42
4.1	Reduction to dimension 1	45

1 Introduction

We study one way in which stability and order can interact in an *NIP* theory. More precisely, we are interested in the situation where stability and order are intertwined. We start by giving some very simple examples illustrating what we mean.

Consider $M_0 \models \text{DLO}$. A type of $S_1(M_0)$ is determined by a cut in M_0 and two types corresponding to different cuts are orthogonal. If we take now M_1 a model of some o-minimal theory, still a 1-type is determined by a cut, but in general, types that correspond to different cuts are not orthogonal. However this is true over indiscernible sequences in the following sense: assume $\langle a_t : t < \omega + \omega \rangle \subset M_1$ is an indiscernible sequence. By *NIP*, the sequences of types $\langle \text{tp}(a_t/M_1) : t < \omega \rangle$ and $\langle \text{tp}(a_{\omega+t}/M_1) : t < \omega \rangle$ converge in $S(M_1)$. Then the two limit types are orthogonal (this follows from dp-minimality, see 2.30). An indiscernible sequence with that property will be called *distal*¹. A theory is distal if all indiscernible sequences are distal. So any o-minimal theory is distal.

Distality for an indiscernible sequence can be considered as an opposite notion to that of total indiscernibility.

¹Thanks to Itay Kaplan for suggesting the name.

Let now M_2 be a model of ACVF (or any other C-minimal structure) and consider an indiscernible sequence $(a_i)_{i < \omega}$ of elements from the valued field sort. Two different behaviors are possible: either the sequence is totally indiscernible, this happens if and only if $\text{val}(a_i - a_j) = \text{val}(a_{i'} - a_{j'})$ for all $i \neq j$, $i' \neq j'$, or the sequence is distal. Again, this will follow from the results in Section 2, but could be proved directly. So M_2 is neither stable nor distal; the two phenomena exist but do not interact in a single indiscernible sequence of points.

Consider now a fourth structure (a ‘colored order’) M_3 in the language $L_3 = \{\leq, E\}$: M_3 is totally ordered by \leq and E defines an equivalence relation, each E class being dense co-dense with respect to \leq . Now an indiscernible sequence of elements from different E classes is neither totally indiscernible nor distal. Given two limit types p_x and q_y of different cuts in such a sequence, the type $p_x \cup q_y$ is consistent with xEy and with $\neg xEy$. Here it is clear that the ‘stable part’ of a type should be its E -class.

The idea behind the work in this paper is that every ordered indiscernible sequence in an *NIP* theory should look like a colored order: there is an order for which different cuts are orthogonal and something stable on top of it which does not see the order (see Section 3).

1.0.1 A word about measures

Keisler measures will be used a little in this work, however the reader not familiar with them can skip all parts referring to measures without harm. For this reason, we will be very brief in recalling some facts about them and refer the reader to [9] and [7]. They however give some understanding of the intuition behind some definitions and results. We explain this now.

A Keisler measure (or simply a measure) is a Borel probability measure on a type space $S_x(A)$. Basic definitions for types (non-forking, invariance, coheir, Morley sequence etc.) generalize naturally to measures (see [9] and [7]). Of interest to us is the notion of *generically stable measure*. A measure is generically stable if it is both definable and finitely satisfiable over some small set. Equivalently, its Morley sequence is totally indiscernible. Such measures are defined and studied by Hrushovski, Pillay and the author in [7]. Furthermore, it is shown in [17] that some general constructions give rise to them, and in this sense they are better behaved than the more natural notion of generically stable type.

This paper can be considered as an attempt to understand where generically

stable measures come from. What stable phenomena do generically stable measures detect? What does the existence of generically stable measures in some particular theory tell us about types? The first test question was: Can we characterize theories which have non-trivial generically stable measures? Here “non-trivial” means “non-smooth”: a measure is smooth if it has a unique extension to any bigger set of parameters. This question is answered in Section 2: a theory has a non-smooth generically stable measure if and only if it is not distal.

The main tool at our disposal to link measures to indiscernible sequences is the construction of an average measure of an indiscernible segment (see [7] Lemma 3.4 or [17] Section 3 for a more elaborate construction). Such a measure is always generically stable. The intuition we suggest is that the ‘order’ component of the sequence is evened out in the average measure and only the ‘stable’ component remains.

1.0.2 Organization of the paper and main results

The paper is organized as follows. The first section contains some basic facts about *NIP* theories and Keisler measures. We give a number of definitions concerning indiscernible sequences and some basic results illustrating how we can manipulate them. Section 2 studies distal theories. They are defined as theories in which every indiscernible sequence is distal, as explained above. We show that this condition can also be seen through invariant types and generically stable measures. The main results can be summarized by the following theorem.

Theorem 1.1. *Let T be *NIP*. Then the following are equivalent:*

- *T is distal,*
- *Any two invariant types that commute are orthogonal,*
- *All generically stable measures are smooth.*

Furthermore, it is enough to check any one of those conditions in dimension 1.

As a consequence, o-minimal theories and the p-adics are distal as are more generally any dp-minimal theory with no generically stable type.

Section 3 can be read almost independently of the previous one: it contains a study of the intermediate case of an *NIP* theory that is neither stable nor

distal. We deal with the problem of understanding to what extent non-distality is witnessed by stable-like interactions between tuples. If M is a $|T|^+$ -saturated model, we define a notion of s-independence denoted $a \downarrow_M^s b$ which is symmetric, is implied by forking-independence and has bounded weight. We use it to show that two commuting types behave with respect to each other like types in a stable theory (we recover some definability and uniqueness of the non-forking extension). The guiding intuition is that of the colored order where elements have a well defined *stable part* (the image in the quotient) and in that case $a \downarrow_M^s b$ means that the stable parts are independent. We do not attempt to give any meaning to the ‘stable part’ of a type in general, and we do not even expect there to be a possible meaning for it. We find that the intuition “s-independence corresponds to independence of stable parts” is useful in understanding those results. Of course, it may turn out some day to be misleading.

As an application of those ideas, we prove the following ‘finite-co-finite theorem’ (Theorem 3.30) and give an application of it to the study of externally definable sets.

Theorem 1.2 (Finite-co-finite theorem). *Assume that T is NIP. Let $I = I_1 + I_2 + I_3$ be indiscernible, I_1 and I_3 being infinite. Assume that $I_1 + I_3$ is A -indiscernible and take $\phi(x; a) \in L(A)$, then the set $B = \{b \in I_2 : \models \phi(b; a)\}$ is finite or co-finite.*

The last section defines a class of theories — called *sharp* — in which (intuitively) the stable part of types is witnessed by generically stable types. More precisely, over a $|T|^+$ -saturated model M , every tuple is s-dominated by the realization of a generically stable type. We give a criterion for sharpness which only involves looking at indiscernible sequences of elements (not tuples). In particular, any dp-minimal theory is sharp.

Our Bible concerning *NIP* theories are Shelah’s papers [11], [12], [10], [14] and [13]. We will however use ideas only from the first two. All the basic insights about indiscernible sequences were taken from there (although the important result on shrinking indiscernible sequences originates in [2]).

In fact, we realized after having done most of this work that the idea of ‘domination’ for indiscernible sequences was already in Shelah’s work: in Section 2 of [12] in a slightly different wording and with a very different purpose. The main additional ingredient in Section 3 is the external characterization of

domination (3.7) which allows us to say something about points outside of the indiscernible sequence and then to generalize to the invariant type setting.

An important property of stable theories sometimes referred to as the *Shelah reflection principle* says roughly that non-trivial relationships between a realization of a type p and some other point are reflected inside realizations of p . Internal concepts (only considering realizations of p) often imply external properties (involving the whole structure). For example regularity implies weight one. There is some evidence now that this principle is already true in *NIP* theories. See [3] for an example (weak stable embeddedness).

In this paper we will use this principle for indiscernible sequences: a property involving only the indiscernible sequence itself or extensions of it usually implies properties of the indiscernible sequence with respect to points outside (the same way total indiscernibility implies that the trace of every definable set is finite or co-finite). See Lemma 2.7 and Proposition 3.7.

Acknowledgment

I would like to thank Elisabeth Bouscaren for having patiently listened to the expositions of many erroneous proofs of results here and for helping finding mistakes. Thanks to Udi Hrushovski for interesting remarks on earlier versions of this work and for his encouragements to pursue it. Thanks also to Tom Scanlon and the anonymous referee for a number of helpful comments.

1.1 Preliminaries

We work with a complete theory T , in a language L . We let \mathfrak{C} denote a monster model of T .

We will often denote sequences of tuples by I, J, \dots . Index sets of families or sequence might be named $\mathcal{I}, \mathcal{J}, \dots$.

If I is an indiscernible sequence and A a set of parameters, let $\lim(I/A)$ let I be the limit type of I over A defined as follows: if $I = (a_t)_{t \in \mathcal{I}}$, and $\phi(x; d) \in L(A)$, then $\phi(x; d) \in \lim(I/A)$ if and only if for some $t_0 \in \mathcal{I}$, $\models \phi(a_t; d)$ holds for all $t \geq t_0$. Recall that a theory is *NIP* if and only if $\lim(I/A)$ is a complete type for every I and A . By $\lim(I)$, we mean the global type $\lim(I/\mathfrak{C})$.

Assumption : Throughout the paper, we assume that the theory T is *NIP*.

Let M be a κ -saturated model (for some $\kappa > |T|$). If $A \subseteq M$, $|A| < \kappa$, then a type $p \in S(M)$ is A -invariant if for $a \models p$ and any tuples $b, b' \in M$,

$b \equiv_A b' \rightarrow ba \equiv b'a$. We will sometimes say simply that p is an invariant type, without specifying A . Note that an invariant type has a natural extension to any larger set $B \supset M$ that we will denote by $p|_B$. We use the same notation to denote the restriction of p to B , when $B \subset M$.

Let \mathcal{I} be a linear order. A Morley sequence indexed by \mathcal{I} of an invariant type p over some $B \supseteq A$ is a sequence $(a_t)_{t \in \mathcal{I}}$ such that $a_t \models p|_{B \cup a_{<t}}$ for every t . All Morley sequences of p over B indexed by \mathcal{I} are B -indiscernible and have the same type over B ; when $B = M$, we will denote that type by $p^{(\mathcal{I})}$.

If p_x and q_y are two types over the κ -saturated model M and p is invariant, we can define the product $p_x \otimes q_y$ as the element of $S_{xy}(M)$ defined as $\text{tp}(a, b/M)$ where $b \models q_y$ and $a \models p_x|_{Mb}$. If q is also an invariant type, then $p_x \otimes q_y$ is invariant. In this case, we can also build the product $q_y \otimes p_x$. When the two products are equal, we say that p and q commute.

Note that \otimes is associative. In particular if p and q commute with r , then r commutes with $p \otimes q$.

Definition 1.3. Two types p_x, q_y over the same domain A are *weakly orthogonal* if $p_x \cup q_y$ defines a complete type in two variables over A .

If $p_x, q_y \in S(M)$ are invariant over $A \subset M$ (M is κ -saturated and $|A| < \kappa$), then we say that p_x and q_y are *orthogonal* if they are weakly orthogonal. This implies that $p|_B$ and $q|_B$ are also weakly orthogonal for any $B \supseteq M$.

Recall the notion of generically stable type from [11] and [9]: an invariant type $p \in S(M)$ is generically stable if it is both definable and finitely satisfiable in some small model $N \subset M$. Equivalently, its Morley sequence is totally indiscernible.

1.1.1 Measures

As we mentioned in the introduction, we will not recall all definitions concerning measures. Instead, we refer the reader to [9] and [7]. The latter paper contains in particular the definition of a generically stable measure. Also the introduction of [17] contains a concise account of the definitions and basic results we will need, but without proofs.

We will need to extend the definition of weakly orthogonal for a type and a measure: if μ_x is a measure over A and p_y a type over the same A , we say that they are *weakly orthogonal* if μ_x has a unique extension to a measure over Ab , where $b \models p_y$.

We also recall the following from [17]: if M is a model, a measure $\mu \in \mathcal{M}_x(M)$ is *smooth* if it has a unique extension to any $N \supset M$. For any formula $\phi(x, d)$, $d \in \mathfrak{C}$, let $\partial_M \phi$ denote the closed subset of $S_x(M)$ consisting of types p such that there are a, a' two realizations of p satisfying $\phi(a, d) \wedge \neg \phi(a', d)$.

Fact 1.4 (Lemma 4.1 of [17]). *The measure $\mu \in \mathcal{M}_x(M)$ is smooth if and only if $\mu(\partial_M \phi) = 0$ for all formulas $\phi(x, d)$, $d \in \mathfrak{C}$.*

1.1.2 Indiscernible sequences and cuts

The notation $I = I_1 + I_2$ means that the sequence I is the concatenation of the sequences I_1 and I_2 : I_1 is an initial segment of I and I_2 the complementary final segment. This operation is associative, and we will also use it to denote the concatenation of three or more sequences. It may be the case that one of the sequences is finite. In particular, when b is a tuple, we may write $I_1 + b + I_2$ to denote $I_1 + \langle b \rangle + I_2$ where $\langle b \rangle$ is the sequence of length 1 whose only member is b .

If $I = I_1 + I_2$, we will say that (I_1, I_2) is a *cut* of I .

By the EM-type (over A) of an indiscernible sequence $I = \langle a_i : i \in \mathcal{I} \rangle$, we mean the family $(p_n)_{n < \omega}$, where $p_n \in S_n(A)$ is the type of $(a_{\sigma(k)})_{k < n}$ for $\sigma : n \rightarrow \mathcal{I}$ any increasing embedding.

We now introduce a number of definitions that will be useful for handling indiscernible sequences.

Definition 1.5 (Cuts). If $J \subset I$ is a convex subsequence, a cut $\mathfrak{c} = (I_1, I_2)$ is said to be *interior* to J if $I_1 \cap J$ and $I_2 \cap J$ are infinite.

A cut is *Dedekind* if both I_1 and I_2^* (I_2 with the order reversed) have infinite cofinality.

If $\mathfrak{c} = (I_1, I_2)$ and $\mathfrak{d} = (J_1, J_2)$ are two cuts of the same sequence I , then we write $\mathfrak{c} \leq \mathfrak{d}$ if $I_1 \subseteq J_1$.

We write $(I'_1, I'_2) \triangleleft (I_1, I_2)$ if I'_1 is an end segment of I_1 and I'_2 an initial segment of I_2 . A *polarized cut* is a pair $(\mathfrak{c}, \varepsilon)$ where \mathfrak{c} is a cut (I_1, I_2) and $\varepsilon \in \{1, 2\}$ is such that I_ε is infinite. We will write the polarized cut \mathfrak{c}^- if $\varepsilon = 1$ and \mathfrak{c}^+ if $\varepsilon = 2$.

Given a polarized cut $\mathfrak{c}^\bullet = ((I_1, I_2), \varepsilon)$ and a set A of parameters, we can define the *limit type* of \mathfrak{c}^\bullet denoted by $\lim(\mathfrak{c}^\bullet/A)$ as the limit type of the sequence I_1 or I_2^* depending on the value of ε .

If a cut \mathfrak{c} has a unique polarization, or if we know both polarizations give the same limit type over A , we will write simply $\lim(\mathfrak{c}/A)$.

If $\mathbf{c} = (I_1, I_2)$ is a cut, we say that the tuple b *fills* the cut \mathbf{c} if $I_1 + b + I_2$ is indiscernible. Similarly, if \bar{b} is a sequence of tuples, we will say that \bar{b} fills c if the concatenation $I_1 + \bar{b} + I_2$ is indiscernible.

The following definition is from [11].

Definition 1.6. Let $\mathbf{c} = (I_1, I_2)$ be a Dedekind cut. A set A *weakly respects* \mathbf{c} if $\lim(\mathbf{c}^+/A) = \lim(\mathbf{c}^-/A)$. It *respects* \mathbf{c} if for every finite $A_0 \subseteq A$, there is I'_1 cofinal in I_1 and I'_2 coinital in I_2 such that $I'_1 + I'_2$ is indiscernible over A_0 .

Note that $\lim(\mathbf{c}^\bullet) = \lim(\mathbf{c}^\bullet/\mathcal{C})$ is an invariant type, in fact finitely satisfiable over the sequence I .

If \mathbf{c}_1 and \mathbf{c}_2 are two distinct polarized cuts in an indiscernible sequence I then $\lim(\mathbf{c}_1)$ and $\lim(\mathbf{c}_2)$ commute: $\lim(\mathbf{c}_1)_x \otimes \lim(\mathbf{c}_2)_y = \lim(\mathbf{c}_2)_y \otimes \lim(\mathbf{c}_1)_x$. More precisely $\phi(x, y) \in \lim(\mathbf{c}_1)_x \otimes \lim(\mathbf{c}_2)_y$ if and only if for some J_1 cofinal in \mathbf{c}_1 and J_2 cofinal in \mathbf{c}_2 , $\phi(a, b)$ holds for $(a, b) \in J_1 \times J_2$.

Definition 1.7 (Polycut). A polycut is a sequence $(\mathbf{c}_i)_{i \in \mathcal{I}}$ of pairwise distinct cuts.

The definitions given for cuts extend naturally to polycuts: a polarized polycut is a family of polarized cuts. If $\mathbf{c} = (\mathbf{c}_i)_{i \in \mathcal{I}}$ is a polarized polycut, then we define $\lim(\mathbf{c}) = \bigotimes_{i \in \mathcal{I}} \lim(\mathbf{c}_i)$. It is a type in variables $(x_i)_{i \in \mathcal{I}}$. A tuple $(a_i)_{i \in \mathcal{I}}$ *fills* \mathbf{c} if the sequence I with all the points a_i added in their respective cut is indiscernible. Note that this is stronger than asking that each a_i fills \mathbf{c}_i .

Definition 1.8 (I -independent). Let I be a dense indiscernible sequence, $\mathbf{c}_1, \dots, \mathbf{c}_n$ pairwise distinct cuts in I and a_1, \dots, a_n filling those cuts, then a_1, \dots, a_n are independent over I (or I -independent) if the tuple (a_1, \dots, a_n) fills the polycut $(\mathbf{c}_1, \dots, \mathbf{c}_n)$.

We will use the notation $a \downarrow_I b$ to mean that a and b are independent over I , *i.e.*, that $I \cup \{a\} \cup \{b\}$ remains indiscernible (where $I \cup \{a\} \cup \{b\}$ is ordered so that a and b fall in their respective cuts). Note that this is a symmetric notion.

The proofs in this paper will involve a lot of constructions with indiscernible sequences. We list here the basic results and ideas we will need for that. We tried to encapsulate in lemmas some constructions that we will use often. However, in some cases, the lemmas will not fit exactly our needs. The reader should therefore bear in mind the principles of those constructions more than the statements themselves. The constructions are grouped in three parts: shrinking, expanding and sliding.

1.2 Shrinking

We start with the very important results concerning shrinking of indiscernibles. We give the statement as in [11, Section 3]. See also [1].

Definition 1.9. A finite convex equivalence relation on \mathcal{I} is an equivalence relation \sim on \mathcal{I} which has finitely many classes, all of which are convex subsets of \mathcal{I} .

Proposition 1.10 (Shrinking indiscernibles). *Let A be any set of parameters and $(a_t)_{t \in \mathcal{I}}$ be an A -indiscernible sequence. Let d be any tuple. Let $\phi(x_d; y_0, \dots, y_{n-1}, z)$ be a formula. There is a finite convex equivalence relation \sim on \mathcal{I} such that given:*

- $t_0 < \dots < t_{n-1}$ in \mathcal{I} ;
- $s_0 < \dots < s_{n-1}$ in \mathcal{I} with $t_k \sim s_k$ for all k ;
- $b \in A^{|t|}$,

we have $\phi(d; a_{t_0}, \dots, a_{t_{n-1}}, b) \leftrightarrow \phi(d; a_{s_0}, \dots, a_{s_{n-1}}, b)$.

Furthermore, there is a coarsest such equivalence relation.

Often we will apply this with $A = \emptyset$, in which case b does not appear.

We elaborate a little bit on this statement. We fix some parameter set A , sequence I , tuple d and formula $\phi(x_d; y_0, \dots, y_{n-1}, z)$ such that I is indiscernible over A . Consider the coarsest equivalence relation \sim satisfying the conclusion of Proposition 1.10.

The relation \sim induces a partition of the sequence \mathcal{I} into convex equivalence classes: $\mathcal{I} = \mathcal{I}_1 + \dots + \mathcal{I}_T$. We define also the corresponding partition of I as $I = I_1 + \dots + I_T$.

The $T-1$ cuts $(I_1 + \dots + I_{k-1}, I_k + \dots + I_T)$, for $k < T$, will be called the cuts *induced by (d, ϕ) on I* (over A). For the purpose of this section, we will denote them by $\text{cut}_I(d, \phi; 0) < \dots < \text{cut}_I(d, \phi; T-1)$. Here A is implicit to simplify the notation. Let also $\mathsf{T}_I(d, \phi) = T$ be the number of such cuts.

Let $\mathcal{F}(n, T)$ be the set of non-decreasing functions from n to T . For any $f \in \mathcal{F}(n, T)$ and $b \in A^{|t|}$, there is a truth value $\varepsilon_{d, \phi; I}(f, b)$ such that $\phi(d; a_{t_0}, \dots, a_{t_{n-1}}, b)$ has truth value $\varepsilon_{d, \phi; I}(f, b)$ for any $t_0 < \dots < t_{n-1}$ with $t_k \in I_{f(k)}$ for all $k < T$.

To summarize, the tuple d the sequence I and the set A being fixed, we have associated, to any formula $\phi(x_d; y_0, \dots, y_{n-1}, z)$ an integer $\mathsf{T}_I(d, \phi)$, cuts $\text{cut}_i(d, \phi; I)$ for $i < \mathsf{T}_I(d, \phi)$ and a function $\varepsilon_I(d, \phi) : \mathcal{F}(n, \mathsf{T}_I(d, \phi)) \times A^{|t|} \rightarrow \{\top, \perp\}$. This data completely describes the type of d over IA .

Lemma 1.11. *Let $I = (a_t)_{t \in \mathcal{I}}$ be A -indiscernible with \mathcal{I} of cofinality at least $|T|^+$, then for any finite tuple d , there is an end segment I' of I that is indiscernible over Ad .*

Proof. Simply take I' to be to the right of all the cuts $\text{cut}_I(d, \phi; i)$. \square

Usually, when we consider the type of a tuple d over an indiscernible sequence I , we are not concerned with the exact type, but only with the number of cuts induced by d on I and their relative position with respect to each other. We now define a notion of similarity between types which makes this precise.

Let d be a tuple and $I = (a_t)_{t \in \mathcal{I}}$ an indiscernible sequence. We define a structure $I_{[d]}$ as follows: its universe is $\{a_t : t \in \mathcal{I}\}$, the language contains a binary relation $<_I$ interpreted as the order on I and for each formula $\phi(x_d; y_0, \dots, y_n) \in L$, a n -ary predicate $R_\phi(y_0, \dots, y_{n-1})$ which holds on $(a_{t_0}, \dots, a_{t_{n-1}})$ if and only if $\models \phi(d; a_{t_0}, \dots, a_{t_{n-1}})$.

Definition 1.12. Let I, J be two indiscernible sequence and d, d' two tuples of the same length. We say that $\text{tp}(d/I)$ and $\text{tp}(d'/J)$ are *similar* if $I_{[d]} \equiv J_{[d']}$.

If I and J are indiscernible over A , we say that the two types are *similar over A* if they are similar, in the expanded language $L(A)$.

Note that in particular, if $\text{tp}(d/I)$ and $\text{tp}(d'/I')$ are similar over A , then $\text{tp}(d/A) = \text{tp}(d'/A)$ and the EM-types of I and J over A are the same.

The structure $I_{[d]}$ is bi-interpretable with the structure having same universe, whose language contains the binary relation $<_I$ and for each cut $\text{cut}_i(d, \phi; I)$ a unary predicate interpreted as the left-piece of the cut. When I and J are densely ordered without endpoints (which will almost always be the case), then $\text{tp}(d/I)$ and $\text{tp}(d'/J)$ are similar over A if and only if for all formula ϕ and ψ as above, the following conditions are satisfied:

- $\text{T}_I(d, \phi) = \text{T}_J(d', \phi)$;
- $\epsilon_I(d, \phi) = \epsilon_J(d', \phi)$;
- for all $i < \text{T}_I(d, \phi)$, the cuts $\text{cut}_I(d, \phi; i)$ and $\text{cut}_J(d', \phi; i)$ are either both of infinite cofinality from the left (resp. right) or both of finite cofinality from the left (resp. right);
- for all $i < \text{T}_I(d, \phi)$ and $j < \text{T}_I(d, \psi)$, we have $\text{cut}_I(d, \phi; i) < \text{cut}_I(d, \psi; j)$ if and only if $\text{cut}_{I'}(d', \phi; i) < \text{cut}_{I'}(d', \psi; j)$;
- there are infinitely many elements in I between the cuts $\text{cut}_I(d, \phi; i)$ and $\text{cut}_I(d', \psi; j)$ if and only if there are infinitely many elements in J between the cuts $\text{cut}_J(d', \phi; i)$ and $\text{cut}_J(d', \psi; j)$.

Lemma 1.13. *Let I be a dense indiscernible sequence over a set A , and d a tuple, then there is $I' \subset I$ of size at most $|T| + |d|$ such that $\text{tp}(d/I')$ and $\text{tp}(d/I)$ are similar over A .*

Proof. This is immediate by Löwenheim-Skolem. \square

1.3 Expanding

Let I be an indiscernible sequence over some set A , and d any tuple. We now study how one can extend I to some bigger sequence I' maintaining the similarity type of $\text{tp}(d/I)$ over A .

First, if I is endless, there is a limit type $\lim(I)$ as defined above. If J realizes a Morley sequence of that type over Ad , then $I + J^*$ is indiscernible, where J^* is the sequence J with the opposite order. Also $\text{tp}(d/I + J^*)$ is similar to $\text{tp}(d/I)$ over A .

Consider now a cut $\mathfrak{c} = (I_1, I_2)$ of I . If I_1 is endless, then we can similarly consider K a Morley sequence of $\lim(I_1)$ over IA . Then $I_1 + K^* + I_2$ is indiscernible and $\text{tp}(d/I_1 + K^* + I_2)$ is similar to $\text{tp}(d/I_1 + I_2)$. If I_2 has no first element, then we can similarly extend by realizing a Morley sequence in $\lim(I_2^*)$. Note that unless the cut \mathfrak{c} is induced by (d, ϕ) on I for some formula ϕ , then $\lim(I_1/IA) = \lim(I_2^*/IA)$.

If we want to extend the sequence I by adding elements in different cuts, we can iterate the above procedure. Note that the order in which we chose the cuts does not matter since the different limit types commute with each other.

We therefore conclude the following lemma.

Lemma 1.14. *Let $I = (a_i)_{i \in \mathcal{I}}$ be an indiscernible sequence over some set A . Assume \mathcal{I} is dense without endpoints. Let d be any tuple and let $\mathcal{J} \supset \mathcal{I}$ be any linearly ordered set extending \mathcal{I} . Then there are tuples $(a_i)_{i \in \mathcal{J} \setminus \mathcal{I}}$ such that the sequence $J = (a_i)_{i \in \mathcal{J}}$ is indiscernible over A and $\text{tp}(d/J)$ is similar to $\text{tp}(d/I)$ over A .*

1.4 Sliding

We are now concerned with the situation where we have A , I and d as above, and we want to produce some d' with the same similarity type as d , but such that the cuts induced by d' are different from those induced by d . We see this as *sliding* the point d along the sequence.

We state the result in a slightly more general form..

Lemma 1.15. *Let I, J be two dense sequences, indiscernible over some set A . Assume they have no endpoints and have the same EM-type over A . Let d be any tuple. For any formula ϕ such that $\text{cut}_I(d, \phi; i)$ is well defined, pick a cut $\mathfrak{d}(\phi; i)$ of J such for any ϕ, ψ, i, j for which this makes sense:*

- *the cuts $\text{cut}_I(d, \phi; i)$ and $\mathfrak{d}(\phi; i)$ are either both of infinite cofinality from the left (resp. right) or both of finite cofinality from the left (resp. right);*
- *we have $\mathfrak{d}(\phi; i) < \mathfrak{d}(\psi; j)$ if and only if $\text{cut}_I(d, \phi; i) < \text{cut}_I(d, \psi; j)$;*
- *there are infinitely many elements in J between the cuts $\mathfrak{d}(\phi; i)$ and $\mathfrak{d}(\psi; j)$ if and only if there are infinitely many elements in I between the cuts $\text{cut}_I(d, \phi; i)$ and $\text{cut}_I(d', \psi; j)$.*

Then there is a point e such that $\text{tp}(e/J)$ is similar to $\text{tp}(d/I)$ over A and $\text{cut}_J(e, \phi; i) = \mathfrak{d}(\phi; i)$ for any ϕ and i .

Proof. This translates into finding e with a prescribed type $p(x)$ over AJ . Let $\theta(x; \bar{m}) \in p(x)$, $\bar{m} \subset J$. Also we may assume that $\theta(x; \bar{m})$ is a conjunction of the form

$$\bigwedge_j \phi_j^{\epsilon_j}(x; \bar{m}, b); \quad b \in A, \bar{m} \in J,$$

where ϵ_j is either 0 or 1 depending on the position of the points in \bar{m} with respect to the cuts $\mathfrak{d}(\phi_j; i)$. We can find an injection $\sigma: \bar{m} \rightarrow I$ such that:

- for every m_0, m_1 in \bar{m} , if $m_0 <_J m_1$, then $\sigma(m_0) <_I \sigma(m_1)$;
- for every index j and $m_0 \in \bar{m}$, the relative position of $\sigma(m_0)$ and the cut $\text{cut}_I(a, \phi_j; i)$ on I is the same as that of m_0 and $\mathfrak{d}(\phi; i)$.

Then σ is a partial isomorphism and $a \models \bigwedge_j \phi_j(x; \sigma(\bar{m}))$. Therefore $\theta(x; \bar{m})$ is consistent and by compactness, $p(x)$ is consistent. \square

Corollary 1.16. *Let I, J be two dense sequences with no endpoints indiscernible over some set A of same EM-type over A . Let a and b be tuples of the same length such that $\text{tp}(a/I)$ and $\text{tp}(b/J)$ are similar over A . Let a' be any tuple. Then there is an indiscernible sequence $J' \supseteq J$ and a tuple b' such that $\text{tp}(bb'/J')$ is similar to $\text{tp}(aa'/I)$ over A .*

Proof. By expanding, we can find a sequence J' extending J such that $\text{tp}(b/J')$ is similar to $\text{tp}(b/J)$ and the sequence J' is indexed by a $|T|^+$ -saturated dense linear order. It is then easy to find cuts $\mathfrak{d}(\phi; i)$ in J' as in the previous lemma corresponding to the cuts $\text{cut}_I(aa', \phi; i)$ in a way compatible with the cuts $\text{cut}_{J'}(b, \phi; i)$ over J' . Lemma 1.15 gives us a tuple $b_0 b'_0$ of same length as aa' such that $\text{tp}(b_0 b'_0/J')$ is similar to $\text{tp}(aa'/I)$. By assumption on the cuts $\mathfrak{d}(\phi; i)$, we have $\text{tp}(b_0/J') = \text{tp}(b/J')$ so by composing by an automorphism over J' , we obtain some b' as required. \square

Corollary 1.17. *Let I, J be two dense sequences with no endpoints indiscernible over A and of same EM-type over A . Let a and b be tuples of the same length such that $\text{tp}(a/I)$ and $\text{tp}(b/J)$ are similar over A . Let $I' \supseteq I$ be indiscernible and let a' be any tuple. Then there is an indiscernible sequence $J' \supseteq J$ and a tuple b' such that $\text{tp}(bb'/J')$ is similar to $\text{tp}(aa'/I')$ over A .*

Proof. Simply apply the previous corollary with a' there equal to $a' \cup (I' \setminus I)$ here. \square

1.5 Weight and dp-minimality

Let $(I_i)_{i < \alpha}$ be a family of indiscernible sequences and A a set of parameters. We say that the sequences $(I_i)_{i < \alpha}$ are *mutually indiscernible* over A if for every $i < \alpha$, the sequence I_i is indiscernible over $A \cup \{I_j : j < \alpha, j \neq i\}$.

The following observations are from [11].

Proposition 1.18. *Let $(I_i)_{i < |T|^+}$ be mutually indiscernible sequences (over some set A) and let d be a tuple of size at most $|T|$. Then there is some $i < |T|^+$ such that I_i is indiscernible over Ad .*

Proof. Assume not, then for every $i < |T|^+$, we can find two tuples \bar{a}_i and \bar{b}_i of increasing elements from I_i and a formula $\phi_i(x, \bar{y})$ such that $d \models \phi_i(x, \bar{a}_i) \wedge \neg \phi_i(x, \bar{b}_i)$. Removing some sequences from the family, we may assume that $\phi_i = \phi$ does not depend on i . By mutual indiscernibility, we have $\text{tp}(a_i/\{I_j : j \neq i\}) = \text{tp}(b_i/\{I_j : j \neq i\})$ for all $i < |T|^+$. It follows that for every $A \subseteq |T|^+$, we can find a tuple d_A such that for all $i < |T|^+$, $d_A \models \phi(x, \bar{a}_i)$ if and only if $i \in A$. This contradicts NIP. \square

Corollary 1.19. *Let M be some κ -saturated model, and let $(p_i)_{i < |T|^+}$ be a family of pairwise commuting invariant types over M . Let $p = \bigotimes_{i < |T|^+} p_i$ and $(a_i)_{i < |T|^+} \models p$. Let also $q \in S(M)$ be any type and $d \models q$. Then there is $i < |T|^+$ such that $(a_i, d) \models p_i \otimes q$.*

Proof. Build a Morley sequence $\langle (a_i^k)_{i < |T|^+} : 0 < k < \omega \rangle$ of p over everything and set $a_i^0 = a_i$ for each i . Commutativity implies that the sequences $(a_i^k)_{k < \omega}$, $i < |T|^+$ are mutually indiscernible. The result then follows by Proposition 1.18. \square

Observe in particular that if q is an invariant type, taking $b \models q|\{a_i : i < |T|^+\}$, we obtain that there is $i < |T|^+$ such that p_i and q commute.

We will occasionally mention dp-minimal theories. They are theories for which the notion of weight suggested by Proposition 1.18 is equal to 1 on 1-types. This notion was introduced by Shelah in [10].

Definition 1.20 (Dp-minimal). A theory T is dp-minimal if it is *NIP* and if for every indiscernible sequence I and element d of the home sort, there is a subdivision $I = I_1 + I_2 + I_3$ into convex sets, where I_2 is either reduced to a point or empty and I_1 and I_3 are both indiscernible over d .

Equivalently, for every two mutually indiscernible sequences I and J and element d , one of I or J is indiscernible over d .

See [16] for the proof of the equivalence and [5] for additional information.

Examples of dp-minimal theories include o-minimal and C-minimal theories and the p-adics.

2 Distal theories

2.1 Indiscernible sequences

We now state the main definition of this paper.

Definition 2.1 (Distal). An indiscernible sequence I is distal if for any dense sequence J of same EM-type as I , and any distinct Dedekind cuts \mathfrak{c}_1 and \mathfrak{c}_2 of J , if a fills \mathfrak{c}_1 and b fills \mathfrak{c}_2 , then $a \downarrow_J b$.

An *NIP* theory T is *distal* if all indiscernible sequences are distal.

Remark 2.2. Equivalently the two types $\lim(\mathfrak{c}_1/J)$ and $\lim(\mathfrak{c}_2/J)$ are weakly orthogonal.

Lemma 2.3. *If I is dense and has two distinct Dedekind cuts \mathfrak{c}_1 and \mathfrak{c}_2 , then it is distal if and only if $\lim(\mathfrak{c}_1/I)$ and $\lim(\mathfrak{c}_2/I)$ are weakly orthogonal (i.e., there is no need for J in the definition).*

Proof. Left to right is obvious. We show the converse. If I is not distal, then there is some dense sequence J of the same EM-type, two distinct Dedekind cuts \mathfrak{d}_1 and \mathfrak{d}_2 of J , some a_1 filling \mathfrak{d}_1 and a_2 filling \mathfrak{d}_2 such that $a_1 \not\downarrow_J a_2$. Let $\phi(a_1, a_2, \bar{m})$ be a formula witnessing that, with $\bar{m} \in I$. Take a countable $J' \subseteq J$ containing \bar{m} such that a_1 and a_2 fill Dedekind cuts of J' . Replacing J by J' , we may assume that J is countable.

Then by expanding, we can find some $J_0 \supseteq J$ and an automorphism σ mapping J_0 onto I and such that the cut \mathfrak{d}_1 (resp. \mathfrak{d}_2) is mapped to \mathfrak{c}_1 (resp. \mathfrak{c}_2) and the types $\text{tp}(a_1, a_2/J)$ and $\text{tp}(a_1, a_2/J_0)$ are similar. Then, the points $\sigma(a_1)$ and $\sigma(a_2)$ fill respectively the cuts \mathfrak{c}_1 and \mathfrak{c}_2 and $\phi(\sigma(a_1), \sigma(a_2), \sigma(\bar{m}))$ holds. Therefore $\sigma(a_1) \perp_I \sigma(a_2)$ and it follows that the two limit types $\lim(\mathfrak{c}_1/I)$ and $\lim(\mathfrak{c}_2/I)$ are not weakly orthogonal. \square

Actually, it will follow from Lemma 2.7 that the hypothesis that I is dense can be removed.

EXAMPLE 2.4. Assume I is an indiscernible sequence, f a definable function such that $f(I)$ is totally indiscernible (non constant), then I is not distal. To see this, take a and b in the definition such that $f(a) = f(b)$. See 2.15 for a more general result.

EXAMPLE 2.5. In DLO, any two 1-types concentrating on different cuts are weakly orthogonal. It is easy then to check that it is a distal theory. We will see (Corollary 2.30) that in fact any o-minimal theory is distal.

Lemma 2.6. Assume I is a dense indiscernible distal sequence, and $\mathfrak{c}_0, \dots, \mathfrak{c}_{n-1}$ are pairwise distinct Dedekind cuts. If for each $i < n$, a_i fills \mathfrak{c}_i then the family $(a_i)_{i < n}$ is I -independent.

Proof. We prove it by induction on n . for $n = 2$, it is Lemma 2.3. Assume it holds for n and consider a family $(\mathfrak{c}_i)_{i < n+1}$ and $(a_i)_{i < n+1}$ as in the hypothesis. Let $I' = I \cup \{a_0\}$ (where a_0 is inserted in the cut \mathfrak{c}_0). Each cut \mathfrak{c}_i naturally induces a cut \mathfrak{c}'_i of I' . By the case $n = 2$, for each $0 < i < n+1$, a_i fills \mathfrak{c}'_i . The sequence I' is also distal, so by induction $(a_i)_{0 < i < n+1}$ is I' -independent. Therefore $(a_i)_{i < n+1}$ is I -independent. \square

Lemma 2.7 (External characterization of distality). *A sequence I is distal if and only if the following property holds: For every set A , tuple b and A -indiscernible sequence $I' = I_1 + I_2$ (I_1 and I_2 without endpoints, $\text{EM-tp}(I') = \text{EM-tp}(I)$), if $I_1 + b + I_2$ is indiscernible, it is A -indiscernible.*

Proof. Assume that I is distal, but the conclusion does not hold. Then there is some $I' = I_1 + I_2$ and formula $\phi(x)$ with parameters from $A \cup I_1 \cup I_2$ which witnesses it. This means $\phi(b)$ holds and there is $(I'_1, I'_2) \triangleleft (I_1, I_2)$ such that $\neg\phi(a)$ holds for $a \in I'_1 \cup I'_2$. Restricting even more if necessary, we may assume that $I'_1 + I'_2$ is indiscernible over the parameters of ϕ . So replacing I' by that

latter sequence, we may assume that all the parameters are from A . Then, we may freely enlarge I' , so assume that it is dense.

As I' is A -indiscernible, for every cut \mathfrak{c} of I' , there is b' filling it such that $\phi(b')$ holds. Fix an increasing sequence $(\mathfrak{c}_k)_{k < \omega}$ of such cuts. For every $k < \omega$, let b_k fill \mathfrak{c}_k such that $\phi(b_k)$ holds. The sequence I' is distal (because I' and I have same EM-type) so by Lemma 2.6, the sequence formed by adding all those points to I' is still indiscernible. Therefore $\phi(x)$ has infinite alternation number, contradicting NIP .

The converse is easy. □

The following technical lemma will be used repeatedly.

Lemma 2.8 (Strong base change). *Let I be an indiscernible sequence and $A \supseteq I$ a set of parameters. Let $(\mathfrak{c}_i)_{i < \alpha}$ be a sequence of pairwise distinct polarized Dedekind cuts in I . For each $i < \alpha$ let d_i fill the cut \mathfrak{c}_i . Then there exist $(d'_i)_{i < \alpha}$ such that $\text{tp}((d'_i)_{i < \alpha}/I) = \text{tp}((d_i)_{i < \alpha}/I)$ and for each $i < \alpha$, $\text{tp}(d'_i/A) = \lim(\mathfrak{c}_i/A)$.*

Proof. Assume the result does not hold. Then by compactness, we may assume that $\alpha = n$ is finite and that there is a formula $\phi(x_0, \dots, x_{n-1}) \in \text{tp}((d_i)_{i < n}/I)$ and formulas $\psi_i(x_i) \in \lim(\mathfrak{c}_i/m)$ for some finite $m \in A^k$ such that $\phi(x_0, \dots, x_{n-1}) \wedge \bigwedge_i \psi_i(x_i)$ is inconsistent. Let I_0 denote the parameters of ϕ , and assume $I_0 \subseteq m$.

Assume for simplicity that $n = 2$ (the proof for $n > 2$ is the same) and without loss each \mathfrak{c}_i is polarized as \mathfrak{c}_i^- . For $i = 0, 1$ take $(J_i, J'_i) \leq \mathfrak{c}_i$ such that ψ_i holds on all elements of J_i and $J_i \cup J'_i$ contains no element of I_0 . Then $J_0 + J'_0$ and $J_1 + J'_1$ are mutually indiscernible over I_0 . So for every two cuts \mathfrak{d}_0 and \mathfrak{d}_1 respectively from $J_0 + J'_0$ and $J_1 + J'_1$, we can find points e_0 and e_1 filling those cuts (even seen as cuts of I) such that $\phi(e_0, e_1)$ holds.

Take two cuts \mathfrak{d}_0 and \mathfrak{d}_1 of I such that they are respectively interior to J_0 and J_1 . Fill \mathfrak{d}_0 by e_0 and \mathfrak{d}_1 by e_1 such that $\phi(e_0, e_1)$ holds. By hypothesis, either $\neg\psi_0(e_0)$ or $\neg\psi_1(e_1)$ holds. Assume $\neg\psi_1(e_1)$ holds. Now forget about e_0 and set $I' = I \cup \{e_1\}$. Then I' is indiscernible and we take it as our new I . Set $J'_0 = J_0$ and let J'_1 be an initial segment of J_1 not containing \mathfrak{d}_1 and make the same construction. We obtain new points (e_0^1, e_1^1) that fill the cuts $\mathfrak{d}_0^1, \mathfrak{d}_1^1$ of J'_0 and J'_1 such that $\neg\psi_0(e_0^1) \vee \neg\psi_1(e_1^1)$ holds. Without loss (as we will iterate infinitely many times) again $\neg\psi_1(e_1^1)$ holds.

Iterate this ω time to obtain a sequence of points e_1^k and cuts \mathfrak{d}_1^k in J_1 such that I with all the points e_1^k added in the cuts \mathfrak{d}_1^k is indiscernible and $\neg\psi_1(e_1^k)$ holds for all n . But $\psi_1(x)$ holds for all $x \in J_1$ so ψ_1 has infinite alternation rank, contradicting NIP . □

Corollary 2.9 (Base change). *The notion of being distal is stable both ways under base change: If I is A -indiscernible, then I is distal in $T(A)$ if and only if it is distal in T .*

Proof. Assume I is distal in T . Notice that the property stated in Lemma 2.7 is preserved under naming parameters (because we can incorporate them in the set A). This implies that I is distal in $T(A)$.

Conversely, assume I is not distal in T . Increase I to some large A -indiscernible sequence $J_1 + J_2 + J_3$ and take a, b such that $J_1 + a + J_2 + J_3$ and $J_1 + J_2 + b + J_3$ are indiscernible, but $J_1 + a + J_2 + b + J_3$ is not. By strong base change, we may assume that a and b realize the limit types over A of the cuts they define. Then $J_1 + a + J_2 + J_3$ and $J_1 + J_2 + b + J_3$ are A -indiscernible, giving a counter-example to distality in $T(A)$. \square

Lemma 2.10. *If T is dp-minimal and I is an indiscernible sequence of elements of the home sort which is not totally-indiscernible, then I is distal.*

Proof. Write $I = (d_i)_{i \in \mathcal{I}}$ and assume that it is not totally indiscernible. Working over some base A if necessary, we may assume that there is a formula $\phi(x, y) \in L(A)$ which orders the sequence I and such that I is indiscernible over A . So we have $\phi(d_i, d_j) \iff i < j$. (Extend the sequence I to some $J_1 + I + J_2$ and take $A = J_1 + J_2$.)

Without loss \mathcal{I} is a dense order and can be written as $\mathcal{I}_1 + \mathcal{I}_2 + \mathcal{I}_3$, the three pieces being infinite without end points. Write $I = I_1 + I_2 + I_3$ in the obvious way. Let a fill the cut $\mathfrak{c}_a = (I_1, I_2 + I_3)$ and b fill $\mathfrak{c}_b = (I_1 + I_2, I_3)$. Assume that a and b contradict distality of I . So there is a formula $\psi(x, y) \in L(AI)$ such that $\psi(a, b)$ holds and witnesses $a \not\perp_I b$. Let $\bar{d} = (d_{i_1}, \dots, d_{i_n})$ be the parameters of ψ coming from I with $i_1 < \dots < i_n$. Let s be such that exactly i_1, \dots, i_s are from \mathcal{I}_1 and t such that exactly i_{s+1}, \dots, i_t are from \mathcal{I}_2 . Let \mathcal{I}'_1 be an end segment of \mathcal{I}_1 above i_s and \mathcal{I}'_3 an initial segment of \mathcal{I}_3 below i_{t+1} .

Let $\bar{d}_1 = (d_{i_1}, \dots, d_{i_s})$ and $\bar{d}_3 = (d_{i_{t+1}}, \dots, d_{i_n})$. Consider the sequence $J = \langle d_i \hat{\bar{d}}_1 \hat{\bar{d}}_3 : i \in \mathcal{I}'_1 + \mathcal{I}_2 \rangle + \langle b \hat{\bar{d}}_1 \hat{\bar{d}}_3 \rangle + \langle d_i \hat{\bar{d}}_1 \hat{\bar{d}}_3 : i \in \mathcal{I}'_3 \rangle$. It is an indiscernible sequence. By dp-minimality applied to J and a , we know that J breaks into $J_1 + J_2 + J_3$, J_2 having at most one element, and such that J_1 and J_3 are indiscernible over a . Considering the formula $\phi(x, a)$, we know that J_1 must be equal to $\langle d_i \hat{\bar{d}}_1 \hat{\bar{d}}_3 : i \in \mathcal{I}'_1 \rangle$. And then J_2 is empty and J_3 is the rest of the sequence. In particular the tuple $b \hat{\bar{d}}_1 \hat{\bar{d}}_3$ lies inside J_3 as do all the parameters of $\psi(x, y)$. As $\psi(a, b)$ holds but $\neg\psi(a, d_i)$ holds for $i \in \mathcal{I}_3$, we get a contradiction to the indiscernibility of J_3 over a . \square

Lemma 2.11. *Let T be distal, I and J are two mutually indiscernible sequences. Let \mathfrak{c} (resp. \mathfrak{d}) be a cut in the interior of I (resp. J). Then $\lim(\mathfrak{c}/IJ)$ and $\lim(\mathfrak{d}/IJ)$ are weakly orthogonal.*

Proof. Write $I = (a_i)_{i \in \mathcal{I}}$ and $J = (b_j)_{j \in \mathcal{J}}$. Assume the conclusion does not hold. Then there are $a \models \lim(\mathfrak{c}/IJ)$ and $b \models \lim(\mathfrak{d}/IJ)$ and a formula $\phi(x, y) \in L(IJ)$ such that $\phi(a, b)$ holds, but $\lim(\mathfrak{c}) \otimes \lim(\mathfrak{d}) \vdash \neg\phi(x, y)$. Let \mathcal{K} be a countable dense linear order without end points. Pick embedding $\tau_1 : \mathcal{K} \rightarrow \mathcal{I}$ and $\tau_2 : \mathcal{K} \rightarrow \mathcal{J}$ such that:

- \mathfrak{c} induces a Dedekind cut on $\tau_1(\mathcal{K})$ and \mathfrak{d} induces a Dedekind cut on $\tau_2(\mathcal{K})$;
- identifying $\tau_1(\mathcal{K})$ and $\tau_2(\mathcal{K})$, those two Dedekind cuts are distinct;
- the parameters of $\phi(x, y)$ belong to $\{a_i : i \in \tau_1(\mathcal{K})\} \cup \{b_j : j \in \tau_2(\mathcal{K})\}$.

Let K be the sequence $\langle a_{\tau_1(t)} \hat{\ } b_{\tau_2(t)} : t \in \mathcal{K} \rangle$. Let \mathfrak{c}' and \mathfrak{d}' denote the two cuts naturally induced by \mathfrak{c} and \mathfrak{d} on K . There are tuples b_* and a_* such that $a \hat{\ } b_*$ fills \mathfrak{c}' and $a_* \hat{\ } b$ fill \mathfrak{d}' . By distality of K , $a \hat{\ } b_* \downarrow_K a_* \hat{\ } b$ and $\phi(a, b)$ holds. This contradicts the assumption. \square

Definition 2.12 (Weakly linked). Let $\langle (a_i, b_i) : i \in \mathcal{I} \rangle$ be an indiscernible sequence of pairs. We say that $(a_i)_{i \in \mathcal{I}}$ and $(b_i)_{i \in \mathcal{I}}$ are weakly linked if for all disjoint subsets \mathcal{I}_1 and \mathcal{I}_2 of \mathcal{I} , $(a_i)_{i \in \mathcal{I}_1}$ and $(b_i)_{i \in \mathcal{I}_2}$ are mutually indiscernible.

Observation 2.13. 1. *If $\langle (a_i, b_i) : i \in \mathcal{I} \rangle$ is A -indiscernible and $(a_i)_{i \in \mathcal{I}}$ and $(b_i)_{i \in \mathcal{I}}$ are mutually indiscernible, then they are mutually indiscernible over A .*

2. *If $\langle (a_i, b_i) : i \in \mathcal{I} \rangle$ is A -indiscernible and $(a_i)_{i \in \mathcal{I}}$ and $(b_i)_{i \in \mathcal{I}}$ are weakly linked, then they are weakly linked over A .*

Lemma 2.14. *Let $\langle (a_i, b_i) : i \in \mathcal{I} \rangle$ be indiscernible.*

1. *If $(a_i)_{i \in \mathcal{I}}$ and $(b_i)_{i \in \mathcal{I}}$ are weakly linked and $(a_i)_{i \in \mathcal{I}}$ is distal, then $(a_i)_{i \in \mathcal{I}}$ and $(b_i)_{i \in \mathcal{I}}$ are mutually indiscernible.*
2. *If $(b_i)_{i \in \mathcal{I}}$ is totally indiscernible, then $(a_i)_{i \in \mathcal{I}}$ and $(b_i)_{i \in \mathcal{I}}$ are weakly linked.*

Proof. (1). Without loss, we may assume that \mathcal{I} is dense. Pick some finite $\mathcal{I}_2 \subset \mathcal{I}$. Then $(a_i)_{i \notin \mathcal{I}_2}$ is indiscernible over $B = (b_i)_{i \in \mathcal{I}_2}$. By applying repeatedly Lemma 2.7, we obtain that $(a_i)_{i \in \mathcal{I}}$ is indiscernible over B . This is enough.

(2). Assume \mathcal{I} is dense and big enough, take $\mathcal{I}_1 \subset \mathcal{I}$ finite and let $A = (a_i)_{i \in \mathcal{I}_1}$. By shrinking of indiscernibles and using total indiscernibility of $(b_i)_{i \in \mathcal{I}}$, there

is $\mathcal{I}_2 \subset \mathcal{I}$ of size at most $|T|$ such that $(b_i)_{i \in \mathcal{I} \setminus \mathcal{I}_2}$ is indiscernible over A . By indiscernibility of $\langle (a_i, b_i) : i \in \mathcal{I} \rangle$, we may take $\mathcal{I}_2 = \mathcal{I}_1$. Therefore $(a_i)_{i \in \mathcal{I}}$ and $(b_i)_{i \in \mathcal{I}}$ are weakly linked. \square

Corollary 2.15. *Let $\langle (a_i, b_i) : i \in \mathcal{I} \rangle$ be an indiscernible sequence. Assume $(a_i)_{i \in \mathcal{I}}$ is totally indiscernible and $(b_i)_{i \in \mathcal{I}}$ is distal, then $(a_i)_{i \in \mathcal{I}}$ and $(b_i)_{i \in \mathcal{I}}$ are mutually indiscernible.*

2.2 Invariant types

We prove here a characterization of distality in terms of invariant types.

If M is a κ -saturated model, by an invariant type over M , we mean a type $p \in S(M)$ invariant over some $A \subset M$, $|A| < \kappa$. If p and q are two invariant types over M , then we can define the products $p_x \otimes q_y$ and $q_y \otimes p_x$ as explained in the introduction. The types p and q *commute* if those two products are equal.

Lemma 2.16. *Assume T is distal. Let M be κ -saturated and let $p, q \in S(M)$ be invariant types. If $p_x \otimes q_y = q_y \otimes p_x$, then p and q are orthogonal.*

Proof. Let $b \models q$ and let $N < M$ a model of size $< \kappa$ such that p and q are N -invariant. Let $I \subset M$ be a Morley sequence of p over N . Let a realize p , and build I' a Morley sequence of p over Mab . The hypothesis implies that $p^{(\omega)}$ and q commute (as \otimes is associative). Thus $b \models q|_{MI'}$ and in particular, $I + I'$ is indiscernible over Nb . By distality, $I + a + I'$ is also Nb -indiscernible. This proves that $tp(a, b/N)$ is determined.

As this is true for any small N over which p and q are invariant, the types p and q are orthogonal. \square

Proposition 2.17. *The theory T is distal if and only if any two global invariant types p and q that commute are orthogonal.*

Proof. Lemma 2.16 gives one implication. Conversely, assume that T is not distal. Then there is a dense indiscernible sequence I , two distinct Dedekind cuts \mathfrak{c}_1 and \mathfrak{c}_2 and a and b filling them such that $a \not\perp_I b$. By Lemma 2.8 (strong base change), we may assume that $I \subset M$, for M a large saturated model, and $a \models \lim(\mathfrak{c}_1^-/M)$, $b \models \lim(\mathfrak{c}_2^-/M)$. Then the types $p = \lim(\mathfrak{c}_1^-/M)$ and $q = \lim(\mathfrak{c}_2^-/M)$ have the required property. \square

Consider $p, q \in S(M)$ and assume only that p is invariant. Then $p_x \otimes q_y$ is well defined, but $q_y \otimes p_x$ does not make sense a priori. We show now how to define $q_y \otimes p_x$.

Let M be κ -saturated and $p \in S(M)$ an A -invariant type for some $A \subset M$ of size $< \kappa$. We define an M -invariant type $p' \in S(\mathfrak{C})$ as follows: Fix a formula $\phi(x; b) \in L(\mathfrak{C})$ and a maximal Morley sequence (a_1, \dots, a_n) of p over A such that $\neg(\phi(a_i; b) \leftrightarrow \phi(a_{i+1}; b))$ holds for all $i < n$ and each a_i is in M . Set $\phi(x; b) \in p'$ if and only if $\models \phi(a_n; b)$. We will call p' the *inverse* of p over M .

Now if $q_y \in S(M)$ is any type, then we define $q_y \otimes p_x$ to be $p'_x \otimes q_y \in S(M)$. Notice that if q was invariant to begin with, then the two definitions of $q_y \otimes p_x$ coincide. Note also that the associativity relation: $p_x \otimes (q_y \otimes r_z) = (p_x \otimes q_y) \otimes r_z$ holds in all possible cases (each product is well defined if and only if at least two of p, q, r are invariant).

The following generalizes Lemma 2.16, the proof is the same, using Lemma 2.19 to build the Morley sequence I of p inside M .

Lemma 2.18. *Assume T is distal. Let M be κ -saturated ($\kappa \geq |T|^+$), $p \in S(M)$ be A -invariant for some A of size $< \kappa$ and $q \in S(M)$ be any type. If $p_x \otimes q_y = q_y \otimes p_x$, then p and q are orthogonal.*

We record the following lemma for future needs.

Lemma 2.19. *Let M be κ -saturated, $\kappa \geq |T|^+$. Let $p, q \in S(M)$, p being A -invariant for some $|A| < \kappa$. Then there is some $B \subset M$, $|B| < \kappa$, such that $A \subseteq B$ and for $b \models q$ and any $a, a' \in M$ such that $a, a' \models p|_B$, we have $\text{tp}(a, b/A) = \text{tp}(a', b/A)$.*

Proof. Fix a formula $\phi(x, y; c) \in L(A)$ and take (a_1, \dots, a_n) in M a maximal Morley sequence of p over A such that $\neg(\phi(a_i, b; c) \leftrightarrow \phi(a_{i+1}, b; c))$ holds for all $i < n$. Then for each $a \in M$, $a \models p|_{Aa_1 \dots a_n}$ we have $\models \phi(a, b; c) \leftrightarrow \phi(a_n, b; c)$.

Take B to contain all the a_i 's obtain by letting $\phi(x, y; c)$ range in $L(A)$. \square

2.3 Generically stable measures

We prove in this section that distal theories are exactly those theories in which generically stable measures are smooth. We consider this as a justification that distality is a meaningful notion. It was proved in [17] that o-minimal theories and the p-adics have this property. This latter result will be generalized in the next section, where we prove that distality can be checked in dimension 1.

We have two tools at our disposal to link indiscernible sequences of tuples to measures. In one direction, starting with an indiscernible sequence of tuples, we can form the average measure. This construction is defined in [7],

extended in [17] and recalled below. In the opposite direction, starting with a generically stable measure μ (or in fact any invariant measure), we can consider the product $\mu^{(\omega)}$ in variables x_1, x_2, \dots . We then want to realize it in some way. We do this by taking smooth extensions; see the proof of Proposition 2.27.

Let $I = (a_t)_{t \in [0,1]}$ be an indiscernible sequence. We can define the *average measure* μ of I as the global measure defined by $\mu(\phi(x)) = \lambda_0(\{t \in [0,1] : a_t \models \phi(x)\})$, where λ_0 is the Lebesgue measure. That measure is generically stable (in fact definable and finitely satisfiable over I).

The *support* of a measure $\mu \in \mathcal{M}(A)$ is the set of weakly-random types for μ , namely the set of types $p \in S(A)$ such that $p \vdash \neg\phi(x)$ for every formula $\phi(x) \in L(A)$ such that $\mu(\phi(x)) = 0$. We will denote it by $S(\mu)$.

Lemma 2.20. *Let μ be the average measure of the indiscernible sequence $I = (a_t)_{t \in [0,1]}$ (over \mathfrak{C}). Then the support $S(\mu)$ of μ is exactly the set of limit types of cuts of I .*

Proof. First, if $\phi(x)$ is satisfied by some $\lim(\mathfrak{c})$, \mathfrak{c} a cut in I , then $\phi(x)$ holds on a subsequence, cofinal in \mathfrak{c} , and therefore has positive measure. Conversely, let $p(x) \in S(\mu)$ (a global type). For each $\phi(x) \in p$, the set $\{t \in [0,1] : a_t \models \phi(x)\}$ is infinite. By compactness of $[0,1]$, there is $r \in [0,1]$ which is in the closure of all of those sets as $\phi(x)$ varies in $L(\mathfrak{C})$. If I is totally indiscernible, then μ is the unique limit type of I , so assume that this is not the case. Then I is ordered by some formula $\psi(x,y) \in L(\mathfrak{C})$. The type p must satisfy either $\psi(x, a_r)$ or $\psi(a_r, x)$. In the first case, p is equal to the limit type to the left of a_r and in the second case, to the limit type to the right of a_r . □

Proposition 2.21 (Smooth measures imply distality). *Let I be an indiscernible sequence indexed by $[0,1]$, and μ be the average measure of I over some model M . Then μ is smooth if and only if I is distal.*

Proof. Assume μ is not smooth and I is distal. Then there exists a formula $\phi(x,a) \in L(\mathfrak{C})$ such that the set of $p \in S(M)$ such that p neither implies $\phi(x,a)$ nor its negation has positive measure (in other words, $p \in \partial\phi$). We know that the support of μ is exactly the limit types of cuts in I . Therefore, one can find ω such cuts $(\mathfrak{c}_i)_{i < \omega}$ in $\partial\phi$. Remove countably many points from I (thus not affecting any limit types) so that the cuts \mathfrak{c}_i become Dedekind.

Restricting to some sub-interval of $[0,1]$, we may assume that $\phi(x,a)$ has constant truth value on I . Without loss, it holds on all members of I . For each

index i , as $\lim(\mathbf{c}_i) \in \partial\phi$, there is b_i filling the cut \mathbf{c}_i over I such that $\neg\phi(b_i, a)$ holds. As I is distal, the sequence formed by adding all the b_i to I is still indiscernible. But then the formula $\phi(x, a)$ has infinite alternation number.

Conversely, assume that I is not distal. If J is an indiscernible sequence, we write J' for the sequence J with the endpoints removed. We can find a partition $I = I_1 + I_2 + I_3$ and points b_1, b_2 such that $I'_1 + b_1 + I'_2 + I'_3$ and $I'_1 + I'_2 + b_2 + I'_3$ are indiscernible, but $I'_1 + b_1 + I'_2 + b_2 + I'_3$ is not. Without loss, assume that I_1 and I_2 have no last element. By strong base change, we may assume that the types of b_1 and b_2 over M are respectively $\lim(I_1)$ and $\lim(I_2)$. There is a formula ϕ , parameters $i_k \subset I_k$ and b'_1 realizing the same type as b_1 over M such that $\phi(i_1, b_1, i_2, b_2, i_3) \wedge \neg\phi(i_1, b'_1, i_2, b_2, i_3)$ holds. Then the border $\partial\phi$ of $\phi(i_1, x, i_2, b_2, i_3)$ contains all limit types of cuts between i_1 and i_2 and has non zero measure. This proves that μ is not smooth. \square

Corollary 2.22. *If all generically stable measures are smooth, then T is distal.*

Before proving the converse, we generalize some earlier lemmas from types to measures. Recall the following fact (which follows for example from Proposition 3.3 of [8]).

Fact 2.23. *Let μ_x be a measure and $(a_i)_{i \in \mathcal{I}}$ an indiscernible sequence. Let $\phi(x; y)$ be a formula and $\epsilon > 0$. Then, for some N , there do not exist $i_1 < \dots < i_N$ such that $|\mu(\phi(x; a_{i_k})) - \mu(\phi(x; a_{i_{k+1}}))| \geq \epsilon$ for all $j = 1 \dots N - 1$.*

By a measure $\mu_{x_1, x_2, \dots}$ being indiscernible, we mean that for any formula $\phi(x_1, \dots, x_n)$ and any increasing map $\tau : \omega \rightarrow \omega$, we have $\mu(\phi(x_1, \dots, x_n)) = \mu(\phi(x_{\tau(1)}, \dots, x_{\tau(n)}))$. We now state the analogue of the previous fact with an indiscernible sequence of measures, which is Corollary 2.12 of [7].

Fact 2.24. *Let $\mu_{x_1, x_2, \dots}$ be indiscernible, and b any tuple. Fix some formula $\phi(x; y)$ and $\epsilon > 0$, then for some N , there do not exist $i_1 < \dots < i_N$ such that $|\mu(\phi(x_{i_j}; b)) - \mu(\phi(x_{i_{j+1}}; b))| \geq \epsilon$ for $j = 1 \dots N - 1$.*

In particular, if $\mu_{x_1, x_2, \dots}$ is totally indiscernible, *i.e.*, remains indiscernible when we permute the variables, then given $\phi(x; y)$, $\epsilon > 0$ and b , for some N there do not exist i_1, \dots, i_N such that $|\mu(\phi(x_{i_j}; b)) - \mu(\phi(x_{i_{j+1}}; b))| \geq \epsilon$ for $j = 1 \dots N - 1$.

If $I = (a_i)_{i \in \mathcal{I}}$ and μ_x is a measure over $\{a_i : i \in \mathcal{I}\}$, we say that I is μ -indiscernible if for all $\phi(x; y_1, \dots, y_n)$, for all $t_1 < \dots < t_n$ and $s_1 < \dots < s_n$ in \mathcal{I} we have $\mu(\phi(x; a_{t_1}, \dots, a_{t_n})) = \mu(\phi(x; a_{s_1}, \dots, a_{s_n}))$.

Lemma 2.25. *Let $I_1 + a + I_2$ be an indiscernible distal sequence, I_1 and I_2 without endpoints. If μ_x is a measure such that $I_1 + I_2$ is μ -indiscernible, then $I_1 + a + I_2$ is also μ -indiscernible.*

Proof. The proof is the same as that of Lemma 2.7 using Fact 2.23. \square

Lemma 2.26. *Let $\mu_{x_1, x_2, \dots}$ be totally indiscernible and let $(b_i)_{i < \omega}$ be a distal indiscernible sequence. Assume that the measure $\eta_{(x_1, y_1), (x_2, y_2), \dots}$ is indiscernible, where η is defined by $\eta(\phi(x_1, x_2, \dots; y_1, y_2, \dots)) = \mu(\phi(x_1, x_2, \dots; b_1, b_2, \dots))$. Then the sequence $(b_i)_{i < \omega}$ is μ -indiscernible.*

Proof. This is the analogue of Corollary 2.15. The same proof goes through. Namely, we first use indiscernibility to increase the index set from ω to a dense order \mathcal{I} . Next, let $\phi(x_{t_1}, \dots, x_{t_n}, b_{t_1}, \dots, b_{t_n}; x_s)$ be a formula, $t_1, \dots, t_n \in \mathcal{I}$ are fixed and $\mathcal{J} \subset \mathcal{I}$ is disjoint for those points. Then using the remark following Fact 2.24 and the indiscernibility of the η , we show that $\eta(\phi(x_{t_1}, \dots, b_{t_1}, \dots; x_s))$ is constant as s varies in \mathcal{J} . From this, we conclude that the sequences are weakly linked, namely for any $\mathcal{I}_1, \mathcal{I}_2$ disjoint subsets of \mathcal{I} , the sequence $\langle b_i : i \in \mathcal{I}_1 \rangle$ is μ' -indiscernible, where μ' is the restriction of μ to the variables $(x_i : i \in \mathcal{I}_2)$.

Finally, we show exactly as in Lemma 2.14 (1), that the sequence $(b_i)_{i < \omega}$ is μ -indiscernible. \square

Proposition 2.27. *If T is distal, then all generically stable measures are smooth.*

Proof. Assume that T is distal and take μ a generically stable measure over some $|T|^+$ -saturated model N . The unique global invariant extension of it will also be denoted by μ . Let a be a tuple. We will show that μ and $\text{tp}(a/N)$ are weakly orthogonal.

Let μ' be an extension of μ to Na . Take a smooth extension μ'' of μ' to some $B \supseteq Na$. Let $(B_i)_{i < \omega}$ be a coheir sequence in $\text{tp}(B/N)$, with $B_0 = B$. The measure μ is definable over B , and for each $i < \omega$, we can consider the measure $\mu_{x_i}^i$ which is defined over B_i the same way μ is defined over B (using the canonical bijection from B to B_i).

Consider the measure $\lambda_{\langle x_i, i < \omega \rangle}$ defined as $\bigotimes_{i < \omega} \mu_{x_i}^i$ (this does not depend on the order of the factors since the μ^i 's are generically stable).

Claim: The measure $\lambda_{x_0, x_1, \dots}$ is totally indiscernible over N , namely for every formula $\phi(x_0, \dots, x_{n-1}) \in L(N)$ and any permutation τ of ω , we have $\lambda(\phi(x_0, \dots, x_{n-1})) = \lambda(\phi(x_{\tau(0)}, \dots, x_{\tau(n-1)}))$.

Proof: Note that $\text{tp}(B_1/B_0N)$ is non-forking over N . In particular $\mu_{x_1}^1|_{B_0N}$ does not fork over N (as it is finitely satisfiable in B_1) so by Proposition 3.3 of [7] $\mu^1|_{B_0N} = \mu|_{B_0N}$. This implies, as μ^0 is invariant over B_0 , that $\mu_{x_0}^0 \otimes \mu_{x_1}^1|_{B_0N} = \mu_{x_0}^0 \otimes \mu_{x_1}|_{B_0N}$. On the other hand, $\mu_{x_0}^0 \otimes \mu_{x_1} = \mu_{x_1} \otimes \mu_{x_0}^0$ and as $\mu^0|_N = \mu|_N$, we have $\mu_{x_1} \otimes \mu_{x_0}^0|_N = \mu_{x_1} \otimes \mu_{x_0}|_N$. Putting it all together, we obtain $\mu_{x_0}^0 \otimes \mu_{x_1}^1|_N = \mu_{x_0} \otimes \mu_{x_1}|_N$.

Iterating this we get, $\lambda|_N = \mu^{(\omega)}|_N$. As μ is generically stable, $\lambda_{x_0, x_1, \dots}$ is totally indiscernible over N .

Now define a measure $\eta_{(x_0, y_0), (x_1, y_1), \dots}$ over N , where y_i is a variable of the same size as B , by $\eta(\phi(x_0, x_1, \dots; y_0, y_1, \dots)) = \lambda(\phi(x_0, x_1, \dots; B_0, B_1, \dots))$. By construction, η is a measure of an indiscernible sequence. Lemma 2.26 yields that for any increasing $\sigma : \omega \rightarrow \omega$, and any $\phi(x_0, x_1, \dots; y_0, y_1, \dots)$,

$$\eta(\phi(x_0, x_1, \dots; y_{\sigma 0}, y_{\sigma 1}, \dots)) = \eta(\phi(x_0, x_1, \dots; y_0, y_1, \dots)).$$

Therefore $\mu^0|_{Na} = \mu^1|_{Na} = \mu|_{Na}$. Thus $\text{tp}(a/N)$ and $\mu|_N$ are weakly orthogonal. This proves that μ is smooth. \square

2.4 Reduction to dimension 1

The goal of this section is to prove the following theorem.

Theorem 2.28. *If all sequences of elements of the home sort are distal, then T is distal.*

We first give an informal (and incomplete) proof using measures. Assume all sequences of elements are distal and consider a generically stable measure μ . Then looking at the proof of Proposition 2.27 we see that μ is weakly orthogonal to all 1-types. Then by induction, adding the points one-by-one, μ is weakly orthogonal to every n -type. One could make this proof rigorous, but it seems to require the fact that no type forks over its base. To avoid this hypothesis and the use of measures, we give a purely combinatorial proof.

So we start with a witness of non-distality of the following form:

- a base set of parameters A , and it what follows we work over A (even when not explicitly mentioned);
- an indiscernible sequence $I = (a_i)_{i \in \mathcal{I}}$ with $\mathcal{I} = (0, 1)$ (the usual interval of \mathbb{R}) for simplicity;

- a tuple $b = (b_j)_{j < n}$, some $l \in (0, 1)$ and tuple a such that:
 - a fills the cut “ l^+ ”: $((a_i : i \leq l), (a_i : i > l))$ of I ,
 - I is b -indiscernible,
 - I with a_l replaced by a is not indiscernible over b .

We make some simplifications. First let $m < n$ be the first integer such that $b' = b_{<m}$ satisfies the requirements in place of b . We can add $b_{<m-1}$ as parameters to the base (by base change, or equivalently we can replace a_i by $a'_i = a_i \wedge b_{<m-1}$) and replace b by b_{m-1} . Therefore, we may assume that $|b| = 1$. Next, adding again some parameters to the base, we may assume that for $i \in \mathcal{I}$, $\text{tp}(a/b) \neq \text{tp}(a_i/b)$.

The goal of the construction that follows is to reverse the situation of a and b , *i.e.*, to construct an indiscernible sequence starting with b that is not distal, the non-distality being witnessed by a (or a conjugate of it).

Step 1: Derived sequence

Let $r = \text{tp}(a, b)$. We construct a new sequence $(a'_i)_{i \in \mathcal{I}}$ such that:

- a'_i fills the cut i^+ of I ;
- $\text{tp}(a'_i, b) = r$ for each i ;
- The sequence $\langle (a_i, a'_i) : i \in \mathcal{I} \rangle$ is b -indiscernible.

This is possible by indiscernability of $(a_i)_{i \in \mathcal{I}}$ over b (by sliding, we may choose the a'_i s filling the cuts and then extract).

Step 2: Constructing an array

Using Lemma 2.8 we can iterate this construction to obtain an array $\langle a_i^n : i \in \mathcal{I}, n < \omega \rangle$ and sequence $\langle b_n : n < \omega \rangle$ such that:

- $a_i^0 = a_i$ for each i ;
- for each $i \in \mathcal{I}$, $0 < n < \omega$, the tuple a_i^n realizes the limit type of the cut i^+ of I over $\langle b_k, a_i^k : i \in \mathcal{I}, k < n \rangle$;
- for each $0 < n < \omega$, $\text{tp}(b_n, (a_i^n)_{i \in \mathcal{I}}/I) = \text{tp}(b, (a'_i)_{i \in \mathcal{I}}/I)$.

Claim: For every $\eta : \mathcal{I}_0 \subset \mathcal{I} \rightarrow \omega$ injective, the sequence $\langle a_i^{\eta(i)} : i \in \mathcal{I}_0 \rangle$ is indiscernible, of same EM-type as I .

Proof. Easy, by construction. □

Expanding and extracting, we may assume that the sequence of rows $\langle b_n + (a_i^n)_{i \in \mathcal{I}} : 0 < n < \omega \rangle$ is indiscernible and that $\langle (a_i^n)_{0 < n < \omega} : i \in \mathcal{I} \rangle$ is indiscernible over the sequence $(b_n)_{n < \omega}$.

Step 3: Conclusion

Claim: The sequences $(b_n)_{n < \omega}$ and $\langle (a_i^n)_{i \in \mathcal{I}} : 0 < n < \omega \rangle$ are weakly linked (Definition 2.12).

Proof. Assume for example that some $\phi(b_n, a_i^k)$ holds for all $i \in \mathcal{I}$ and any (k, n) such that $k < n$. Take n very large and take η as in the first claim such that the truth value of “ $\eta(i) < n$ ” alternates more times than the alternation number of ϕ . Then we see that $\phi(b_n, a_i^k)$ must hold also for $k > n$ (otherwise $\phi(b_n, y)$ would alternate too much on the sequence $(a_i^{\eta(i)})$). We can do something similar if the formula ϕ has extra parameters from the b_n ’s or a_i^n ’s, thus it follows that the sequences are weakly linked. \square

Choose an increasing map $\eta : \omega \rightarrow \mathcal{I}$, then the sequences $(b_n)_{n < \omega}$ and $(a_{\eta(n)}^n)_{n < \omega}$ are weakly linked but not mutually indiscernible. This contradicts Lemma 2.14 and finishes the proof of Theorem 2.28.

Corollary 2.29. *If all generically stable measures in dimension 1 are smooth, then all generically stable measures are smooth.*

This generalizes results of [17] where this was proved under additional assumptions.

Corollary 2.30. *If T is dp-minimal and has no generically stable type (in M), then it is distal. In particular o-minimal theories and the p -adics are distal.*

Proof. Recall from 2.10 that in a dp-minimal theory, any indiscernible sequence of elements is either distal or totally indiscernible. \square

Appendix: strong honest definitions

In a later work [4] with Artem Chernikov, we give yet another characterization of distal theories, which is probably the easiest one to use. In particular, one can obtain with it a much shorter proof of the fact that generically stable measures are smooth. We give only the statement here and refer the reader to [4] for more details.

Theorem 2.31. *A theory T is distal if and only if the following holds:*

For any $\phi(x, y)$ there is $\theta(x, z)$ such that: for any finite set C and tuple a , there is $b \in C$ such that $\models \theta(a, b)$ and $\theta(x, b) \vdash \text{tp}_\phi(a/C)$.

3 Domination in non-distal theories

We have now two extreme notions for indiscernible sequences: distality and total indiscernibility. We want to understand the intermediate case. In particular, we want to show that non-distality is witnessed by stable-like phenomena. This part is essentially independent of the previous one but is of course motivated by it. We first concentrate on indiscernible sequences, and then adapt the results to invariant types. A last subsection gives an application to externally definable sets.

The reader might find it useful to have in mind the example of a colored order as defined in the introduction while reading this section.

We will sometimes work with saturated indiscernible sequences, as defined below.

Definition 3.1 (Saturated sequence). An indiscernible sequence of α -tuples is saturated if it is indexed by an $(|T| + |\alpha|)^+$ -saturated dense linear order without end points.

In this section, all cuts are implicitly assumed to be Dedekind (*i.e.*, of infinite cofinality from both sides).

If \bar{a} fills a cut \mathfrak{c} of I , an extension $J \supseteq I$ is *compatible with \bar{a}* if \bar{a} also fills a cut of J .

We fix a global A -invariant type $p \in S_\alpha(\mathfrak{C})$, for some small parameter set A . The indiscernible sequences we will consider will be Morley sequences of p . This is not a real restriction since every indiscernible sequence is a Morley sequence of some invariant type.

The following is the main definition of this section.

Definition 3.2 (Domination). Let I be a dense indiscernible Morley sequence of p over A , $a \models p|_A$ and \mathfrak{c} a cut of I filled by a dense sequence $\bar{a}_* = \langle a_t : t \in \mathcal{I} \rangle$ of α -tuples. We say that \bar{a}_* dominates a over (I, A) if: For every cut \mathfrak{d} of I distinct from \mathfrak{c} , and \bar{b} a dense sequence filling \mathfrak{d} , we have in the sense of $T(A)$:

$$\bar{b} \downarrow_I \bar{a}_* \Rightarrow \bar{b} \downarrow_I a.$$

We say that \bar{a}_* *strongly dominates* a over (I, A) if for every $I \subseteq J$ compatible with \bar{a}_* over A and such that $a \models p|_{AJ}$, \bar{a}_* dominates \bar{a} over J .

We use the notation $\bar{b} \downarrow_I a$ introduced after Definition 1.8 which, in this situation, means $a \models p|_{I\bar{b}}$.

EXAMPLE 3.3. Let T be the theory of colored orders, as defined in the introduction. Let p be an A -invariant type of an element of a new color. Let $I + a$ be a Morley sequence of p over A . Let \mathfrak{c} be a cut in I . If a_* fills \mathfrak{c} , then a_* dominates a over (I, A) if and only if a and a_* have the same color.

Lemma 3.4. *The fact that \bar{a}_* strongly dominates a over (I, A) only depends on the similarity class of $\text{tp}(a, \bar{a}_*/I)$ over A .*

Proof. The statement means that if J is a dense indiscernible sequence, \bar{b}_* and b are tuples such that $\text{tp}(b, \bar{b}_*/J)$ is similar to $\text{tp}(a, \bar{a}_*/I)$ over A , then \bar{b}_* strongly dominates b over (J, A) if and only if \bar{a}_* strongly dominates a over (I, A) . Take such \bar{b}_* , b and J . Assume that $\text{tp}(\bar{b}_*, b/J)$ is similar to $\text{tp}(\bar{a}_*, a/I)$ over A . In particular, J and I have same EM-type over A , so J is also a Morley sequence of p over A . It also follows that $b \models p|_{JA}$ so it makes sense to ask for domination.

Assume that \bar{b}_* does not strongly dominate b over (J, A) . Then we can find a dense sequence $J' \supseteq J$ compatible with \bar{b}_* such that $b \models p|_{J'A}$, some cut \mathfrak{d} of J' and sequence \bar{b}' filling \mathfrak{d} such that $\bar{b}' \downarrow_{J'} \bar{b}_*$, but $\bar{b}' \not\downarrow_{J'} b$ (all over A). By Corollary 1.17 (sliding), we may find $I' \supseteq I$ and \bar{a}' such that $\text{tp}(\bar{b}', \bar{b}_*, b/J')$ is similar to $\text{tp}(\bar{a}', \bar{a}_*, a/I')$ over A . This implies the following facts:

- I' is compatible with \bar{a}_* and $a \models p|_{I'A}$;
- \bar{a}' fills a cut of I' distant from the cut of \bar{a}_* ;
- $\bar{a}' \downarrow_{I'} \bar{a}_*$ and $\bar{a}' \not\downarrow_{I'} a$.

Therefore \bar{a}_* does not strongly dominate a over (I, A) . □

Lemma 3.5. *If \bar{a}_* strongly dominates a over (I, A) , then there is a subsequence $I' \subseteq I$ of size at most $|T| + |\alpha|$ such that \bar{a}_* strongly dominates a over (I', A) .*

Proof. This follows from the previous lemma and Lemma 1.13 (shrinking). □

Proposition 3.6. *Let I be a dense Morley sequence of p over A and $a \models p|_{AI}$, \mathfrak{c} a cut of I then there is a sequence of α -tuples \bar{a}_* of length at most $|T| + |\alpha|$ such that \bar{a}_* fills \mathfrak{c} and \bar{a}_* strongly dominates a over (I, A) .*

Proof. Recall the notation $\mathsf{T}_I(a, \phi)$ from Section 1.2. If $J \subseteq J'$ are two sequences, indiscernible over A , then for any formula ϕ for which this is well defined, we have: $\mathsf{T}_J(a, \phi) \leq \mathsf{T}_{J'}(a, \phi)$. We will write $J \triangleleft J'$ if for some ϕ , this inequality is strict.

Let \mathfrak{J} be the class of indiscernible sequences J such that one can find dense sequences J_1 and J_2 satisfying:

- $J_1 + J + J_2$ is a Morley sequence of p over A ;
- $a \models p|_{AJ_1J_2}$.

If we have a family $(I_i)_{i < \lambda}$ of indiscernible sequences such that $I_i \subseteq I_j$ and $I_i \triangleleft I_j$ hold for all $i < j$, then taking I_λ to be $\bigcup_{i < \lambda} I_i$, we have $I_i \triangleleft I_\lambda$ for all i . Notice in addition that if each I_i belongs to \mathfrak{J} , then it is also the case for I_λ (we can find J_1 and J_2 by compactness). As the numbers $\mathsf{T}_{J'}(a, \phi)$ are finite, it follows that we can find some sequence J in the class \mathfrak{J} such that there is no $J' \supset J$ in this class with $J \triangleleft J'$. By shrinking, we may assume that J is of size $|T| + |\alpha|$. Take J_1 and J_2 as in the definition of \mathfrak{J} . Write $\mathfrak{c} = (I_1, I_2)$. Without loss, J_1 and J_2 have same order types as I_1 and I_2 respectively. Composing by an automorphism over Aa , we may assume that $J_1 = I_1$ and $J_2 = I_2$. Then J fits in the cut \mathfrak{c} . Set $\bar{a}_* = J$.

Assume that \bar{a}_* does not strongly dominate a over (I, A) . Then there is a dense sequence $I' \supseteq I$ a cut \mathfrak{d} of I' and a sequence \bar{b} filling \mathfrak{d} such that:

- \bar{a}_* fills a cut \mathfrak{c}' of I' (over A);
- $a \models p|_{AI'}$;
- $\bar{b} \perp_{I'} \bar{a}_*$, and $\bar{b} \not\perp_{I'} a$.

The sequence $K = I' \cup \bar{a}_* \cup \bar{b}$ (where \bar{a}_* and \bar{b} are placed in their respective cuts) belongs to \mathfrak{J} . Also $\bar{b} \perp_{I'} a$ implies that $\bar{a}_* \triangleleft K$. This contradicts maximality of \bar{a}_* and proves that \bar{a}_* strongly dominates a over (I, A) . \square

3.0.1 External characterization and base change

Similarly to what we did in the distal case, we give an external characterization of domination.

Proposition 3.7 (External characterization of domination). *Let I be a dense Morley sequence of p over A , $a \models p_{AI}$. Let \bar{a}_* fill a cut \mathfrak{c} of I over A such that \bar{a}_* strongly dominates a over (I, A) . Let also $d \in \mathfrak{C}$. Assume:*

- \square *There is a partition $I = J_1 + J_2 + J_3 + J_4$ such that J_2 and J_4 are infinite, \mathfrak{c} in interior to J_2 , $J_2 \cup \{\bar{a}_*\}$ is indiscernible over $Ad + J_1 + J_3 + J_4$ and J_4 is a Morley sequence of p over $Ad + J_1 + J_2 + J_3$.*

Then $a \models p|_{AId}$.

Proof. Let $I, a, \bar{a}_*, d, J_1, \dots, J_4$ as in the statement of the proposition. We may freely enlarge the sequence J_2 , so we may assume that it is saturated (for example, add realizations of limit types of cuts in J_2 over everything. This maintains the hypothesis).

Assume a does not realize p over AId . Then there is some finite $\bar{i} \subset I$ and a formula $\phi(y, \bar{z}; x) \in L(A)$ such that $\models \phi(d, \bar{i}; a)$, but $p \not\models \phi(d, \bar{i}; x)$. Incorporating \bar{i} in d and changing the partition so that $J_2 \cup J_4$ contains no point from \bar{i} , we may assume that $\bar{i} = \emptyset$. Pick a sequence of cuts of J_2 $\mathfrak{c}_0 < \mathfrak{c}_1 < \dots$. Let $\langle \bar{a}_*^k : k < \omega \rangle$ fill the polycut $\langle \mathfrak{c}_k : k < \omega \rangle$ over $Ad \cup \{J_l : l \neq 2\}$, where each \bar{a}_*^k is a sequence of same order type as \bar{a}_* . Let I' denote the sequence I with the points $\bar{a}_*^k, k > 0$, placed in their respective cuts.

Then $\text{tp}(\bar{a}_*^0, d/I')$ is similar to $\text{tp}(\bar{a}_*, d/I)$. By sliding (Corollary 1.16; note that our sequence is already large enough, so we do not need to increase it), we find a_0 such that: $a_0 \models p|_{AI'}$, $\phi(d; a_0)$ holds and \bar{a}_*^0 strongly dominates a_0 over (I', A) .

Let K_1 realize an infinite Morley sequence of p over everything considered so far. Let $I_1 = I \cup \{\bar{a}_*^k : k > 1\} + K_1$ (where the tuples \bar{a}_*^k are placed in their respective cuts). As above, we may find $a_1 \models p|_{AI_1}$ such that \bar{a}_*^1 strongly dominates a_1 over (I_1, A) and $\phi(d; a_1)$ holds. Now as $a_0 \downarrow_{I_1} \bar{a}_*^1$, by the domination assumption we have $a_0 \downarrow_{I_1} a_1$. We iterate this construction building an indiscernible sequence $I_\omega = I + K_1 + K_2 + \dots$ and points $\langle a_k : k < \omega \rangle$ filling the cuts between the K_i 's and independent over I_ω such that $\phi(d; a_k)$ holds for each k . As by assumption $\neg\phi(d; x)$ holds for every $x \in I_\omega$, ϕ has infinite alternation rank, contradicting *NIP*. \square

Proposition 3.8 (Base change). *Let p be A invariant and $A \subset B$. If I is a dense Morley sequence of p over B , $a \models p|_{BI}$ and \bar{a}_* fills a cut of I in the sense of $T(B)$, then if \bar{a}_* strongly dominates a over (I, A) it does so over (I, B) .*

Proof. Assume that \bar{a}_* fills a cut \mathfrak{c} of I in the sense of $T(B)$ and dominates a over (I, A) . Then let \bar{d} fill a cut \mathfrak{c}' of I over B with \mathfrak{c}' distinct from \mathfrak{c} . Assume that $\bar{d} \downarrow_I \bar{a}_*$ over B . Then \boxplus holds with d there replaced by $\bar{d}B$. By domination over (I, A) and the previous proposition, $a \models p|_{I \cup \bar{d}B}$. This proves that \bar{a}_* dominates a over (I, B) . This remains true if we first increase I so \bar{a}_* strongly dominates a over (I, B) . \square

3.1 Domination for types

We now have all we need to state domination results for types over $|T|^+$ -saturated models, instead of cuts in indiscernible sequences.

We work over a fixed κ -saturated model M . By an *invariant type* we mean here a type over M , invariant over some $A \subset M$ of size less than κ .

For the following definition, recall the construction of $p_x \otimes q_y$ when q is invariant (Lemma 2.19 and the paragraph following it).

Definition 3.9 (Distant). Let $p, q \in S(M)$ be two types, assume that at least one of them is invariant, then we say that p and q are distant if they commute: $p_x \otimes q_y = q_y \otimes p_x$ ². If $a, b \in \mathfrak{C}$, we will say that a and b are distant over M if $\text{tp}(a/M)$ and $\text{tp}(b/M)$ are.

Keep in mind that the notion “ a and b are distant over M ” only depends on $\text{tp}(a/M) \cup \text{tp}(b/M)$ and does not say anything more about $\text{tp}(a, b/M)$. In particular, in a stable theory, any a is distant from itself. So distant should not be confused with independent as defined now.

Definition 3.10 (Independent). Given two distant types $p, q \in S(M)$ and $a \models p$, $b \models q$ we say that a and b are independent over M if $\text{tp}(a, b/M) = p \otimes q$. We write $a \downarrow_M b$. This is a symmetric relation.

Definition 3.11 (S-domination). Let $p \in S(M)$ be any type, $a \models p$. A tuple b s -dominates a over M if:

- ⊖ For every invariant type $r \in S(M)$ distant from p and $\text{tp}(b/M)$, and $d \models r$, if $d \downarrow_M b$, then $d \downarrow_M a$.

The reader might be concerned by the fact that this definition depends on the choice of κ (taking a smaller κ we have less invariant types to check). However, we will see later that we get an equivalent definition if we add in ⊖ the condition that r is invariant over a subset of size \aleph_0 .

EXAMPLE 3.12. Taking again the example of a colored order, if p and q are two invariant types (of tuples), $\bar{a} \models p$ and $\bar{b} \models q$, then \bar{b} s -dominates \bar{a} over M if and only if, for every point a_0 in $\text{range}(\bar{a})$, there is a point b_0 in $\text{range}(\bar{b}) \cup M$ of the same color.

²Recall the definition of commuting for non-invariant given after Lemma 2.19

3.1.1 The moving-away lemma

Lemma 3.13. *Let $p \in S(M)$ be any type, and $a \models p$. Then there is some a_* s -dominating a over M and furthermore a_* realizes some invariant type over M .*

Proof. This is similar to Proposition 3.6. Start with some a_* realizing an invariant type. If it does not dominate a , there is an invariant type r distant from a_* and a over M and $b \models r|Ma_*$ such that $b \not\perp_M a$. Replace a_* by a_*b and iterate. By Corollary 1.19, this construction must stop after less than $(|T| + |a|)^+$ steps. \square

For applications we will also need to show that we can find such a dominating tuple distant from any given type.

Lemma 3.14. *Let $I \subset M$ be a dense indiscernible sequence of α -tuples and $(I_i)_{i < \lambda}$ a family of distinct initial segments of I , with $\lambda \geq (|T| + |\alpha|)^+$. For $i < \alpha$, let $p_i = \lim(I_i/M)$. Then given a type $q \in S(M)$, there is $i < \lambda$ such that p_i is distant from q .*

Proof. Observe that the types p_i pairwise commute. Then use Corollary 1.19 (and the remark after it). \square

Lemma 3.15. *Let $p, q \in S(M)$, be types of α -tuples ($|\alpha| < \kappa$) with p invariant over some small A . Let $a \models p$. Then there is $r \in S(M)$ invariant over some B of size \aleph_0 , distant from p and q and $\bar{b} \models r$ such that $|\bar{b}| \leq |T| + |\alpha|$ and \bar{b} s -dominates a over M .*

Proof. By Proposition 3.6 (and Lemma 3.5) we can find I'_0 a dense Morley sequence of p over A of size $|T| + |\alpha|$ and \bar{a}'_* such that $a \models p|AI'_0$, \bar{a}'_* fills a cut \mathfrak{c} of I'_0 and \bar{a}'_* strongly dominates a over (I'_0, A) . Let \bar{b}' be the sequence $I'_0 \cup \bar{a}'_*$ where \bar{a}'_* is placed in its cut.

Let $I \subset M$ be a saturated Morley sequence of p over A , let \mathfrak{c} be a polarized cut of I of cofinality \aleph_0 such that $\lim(\mathfrak{c})$ is distant from q and p (using Lemma 3.14). We may find some $\bar{b} \equiv_{Aa} \bar{b}'$ such that \bar{b} fills the cut \mathfrak{c} of I . Let also I_0, \bar{a}_* be such that $(\bar{b}, I_0, \bar{a}_*) \equiv (\bar{b}', I'_0, \bar{a}'_*)$. So $\bar{b} = I_0 \cup \bar{a}_*$.

Let I_∞ realize an infinite Morley sequence of p over everything. The strong base change lemma (2.8) works equally well if instead of considering points d_i filling the cuts \mathfrak{c}_i , we take sequences \bar{d}_i . We apply this modified version with M as set of parameters, $I + I_\infty$ as indiscernible sequence, $\bar{d}_0 = \bar{b}$ and $\bar{d}_1 = a$. We conclude that we may assume that \bar{b} is a Morley sequence of $\lim(\mathfrak{c})$ over M .

Set $r = \text{tp}(\bar{b}/M)$ and let $B \subset M$ be of size \aleph_0 such that r is B -invariant. Note that r is a power of $\text{lim}(\mathfrak{c})$, so it also commutes with p and q .

Let d realize any invariant type $s \in S(M)$ distant from p and r . Assume that $d \downarrow_M \bar{b}$. Let $C \subset M$ be a subset of size $< \kappa$ such that p, s and r are invariant over C . Let $I' \subset M$ be a Morley sequence of p over C indexed by some dense order \mathcal{I} . Then $d\bar{b}$ realizes $s \otimes r$ over CI' (indeed over M). As p is distant from both r and s , by associativity of \otimes , $p^{(\mathcal{I})}$ commutes with $s \otimes r$. Therefore, I' realizes $p^{(\mathcal{I})}$ over $Cd\bar{b}$. Similarly, \bar{b} realizes r over $CI'd$, and in particular, \bar{b} is indiscernible over $CI'd$.

Furthermore, as $I' \subset M$, \bar{b} realizes r over CI' . As r commutes with p , I' realizes $p^{(\mathcal{I})}$ over $C\bar{b}$, a fortiori over $A\bar{b}$. But \bar{b} is a Morley sequence of p over A . Therefore $\bar{b} + I'$ is a Morley sequence of p over A .

The hypothesis of Proposition 3.7 are satisfied with $J_1 = J_3 = \emptyset$, $J_2 = I_0$, $J_4 = I'$ and d there equal to Cd . We conclude that $a \models p|Cd$. As this is true for every small C , d and a are independent over M . This proves that \bar{b} s -dominates a over M . \square

Remark 3.16. The tuple \bar{b} constructed in the previous lemma has the following additional property:

(D) For every $d \in \mathfrak{C}$ such that $\text{tp}(d/M\bar{b})$ does not fork over M , and such that $\text{tp}(\bar{b}d/M)$ commutes with p , we have $a \downarrow_M d$.

This assumption is satisfied in particular when d is distant from a and \bar{b} , and $\bar{b} \downarrow_M d$ (although d might not realize an invariant type).

Proof. We indicate how to modify the proof above. First, we take C such that p and r are invariant over C . Next take C_1 , $C \subseteq C_1 \subset M$, such that for any $J, J' \subset M$ Morley sequences of p over C_1 indexed by ω , we have $\text{tp}(J/C\bar{b}d) = \text{tp}(J'/C\bar{b}d)$. This is possible using Lemma 2.19. Build I' as a Morley sequence of p over C_1 . By definition of commuting, I' is a Morley sequence of p over $C\bar{b}d$. Also because $\text{tp}(d/M\bar{b})$ does not fork over M , \bar{b} is indiscernible over Md . Finally, the proof that $\bar{b} + I'$ is a Morley sequence of p over A does not change. So as above, we may apply Proposition 3.7 to conclude that d and a are independent over M . \square

Corollary 3.17. *Let $p, q \in S(M)$ be any two types of α -tuples ($|\alpha| < \kappa$) and let $a \models p$. Then there is a_* a tuple of length $\leq |T| + |\alpha|$, distant from q over M and such that a_* s -dominates a over M . Furthermore, we may assume that $\text{tp}(a_*/M)$ is invariant over a subset of size \aleph_0 .*

Proof. By Lemma 3.13, there is some a_{**} s-dominating a over M and realizing some invariant type. By Lemma 3.15, there is a tuple a_* s-dominating a_{**} over M with the required size, whose type over M is invariant over a subset of size \aleph_0 and distant from q .

We check that a_* s-dominates a over M . Let $r \in S(M)$ be an invariant type distant from a_* and a . Let $b \models r$ with $b \downarrow_M a_*$. By Lemma 3.15, there is b_* s-dominating b and distant from $q = \text{tp}(\hat{a} \hat{a}_* \hat{a}_{**}/M)$. Furthermore assume that b_* satisfies property (D). Composing by an automorphism over Mb , we may further assume that $b_* \downarrow_M a_*$. Then as a_* s-dominates a_{**} over M , we have $b_* \downarrow_M a_{**}$ and as a_{**} s-dominates a over M , $b_* \downarrow_M a$. By property (D) this implies $b \downarrow_M a$. \square

Lemma 3.18 (Transitivity of s-domination). *Let $a \in \mathfrak{C}$ and let a_* s-dominate a over M . Let also a_{**} s-dominate a_* over M . Then a_{**} s-dominates a over M .*

Proof. Let $d \in \mathfrak{C}$ be distant from a and a_{**} with $d \downarrow_M a_{**}$. By Corollary 3.17, let d_* s-dominate d over M and distant from $\hat{a} \hat{a}_* \hat{a}_{**}$. Composing by an automorphism over Md , we may assume that $d_* \downarrow_M a_{**}$. Then we have $d_* \downarrow_M a_*$ and $d_* \downarrow_M a$ and finally $d \downarrow_M a$. \square

EXAMPLE 3.19. *If $p \in S(M)$ is generically stable, and $a \models p$, then a is s-dominated by itself. In the opposite situation, if p is invariant and its Morley sequence is distal, then a is s-dominated by the empty set.*

3.1.2 S-independence

Definition 3.20 (S-independence). Let p, q be any types over M , let $a \models p$ and $b \models q$. We say that a and b are *s-independent* over M and write $a \downarrow_M^s b$ if there is a tuple a_* realizing an invariant type, s-dominating a and distant from b such that $a_* \downarrow_M b$.

Note that if a and b are distant, then $a \downarrow_M^s b$ if and only if $a \downarrow_M b$.

Proposition 3.21 (Existence). *Let $p, q \in S(M)$ be any two types and $a \models p$. Then there is $b \models q$ such that $a \downarrow_M^s b$.*

Proof. Let a_* be s-dominating a such that a_* realizes some invariant type p_* distant from q . Take b such that $\text{tp}(a_*, b/M) = p_* \otimes q$. Then by definition $a \downarrow_M^s b$. \square

Proposition 3.22 (Symmetry of s-independence). *S-independence is symmetric: if a and b are two tuples, then $a \downarrow_M^s b$ if and only if $b \downarrow_M^s a$ if and only if there are a_* , b_* s-dominating a and b respectively, distant from each other such that $a_* \downarrow_M b_*$.*

Proof. It is enough to prove the last equivalence. To see right to left, let a_{**} s-dominate a_* and be distant from b_* and b over M . Assume also that $a_{**} \downarrow_M b_*$, then by Lemma 3.18, a_{**} s-dominates a over M . As it is independent from b_* over M , we have $a_{**} \downarrow_M b$ as required.

Conversely, assume that $a \downarrow_M^s b$. Let a_* be a tuple s-dominating a , realizing an invariant type over M , and distant from b such that $b \downarrow_M a_*$. We can find a tuple b'_* s-dominating b distant from a, a_* and b . As $a_* \downarrow_M b$, there is $b_* \equiv_{Mb} b'_*$ such that $a_* \downarrow_M b_*$. \square

Proposition 3.23 (Weight is bounded). *Let $(b_i)_{i < |T|^+}$ be a sequence of tuples such that $b_i \downarrow_M^s b_{<i}$ for each i , and let $a \in \mathfrak{C}$. Then there is $i < |T|^+$ such that $a \downarrow_M^s b_i$.*

Proof. By Lemma 3.17, we can find a family $(b_i^*)_{i < |T|^+}$ such that: For each $i < |T|^+$, b_i^* realizes an invariant type r_i distant from $q := \text{tp}(a/M)$ and r_j , $j \neq i$, b_i^* s-dominates b_i over M and $b_i^* \downarrow_M b_{<i}^*$. By Corollary 1.19, there is $i < |T|^+$ such that $\text{tp}(b_i^*, a/M) = r_i \otimes q$. By definition, $a \downarrow_M^s b_i$. \square

The following special case of this proposition makes no reference to s-domination.

Corollary 3.24. *Let $q \in S(M)$ be A -invariant and, for $i < |T|^+$, let $p_i \in S(M)$ be an invariant type. Assume that p_i commutes with q , for each i . Let $(b_i) \models \bigotimes p_i$ and $a \models q$. Then there is $i < |T|^+$ such that $\text{tp}(b_i, a/N) = p_i \otimes q$.*

Corollary 3.25. *Let $a, b \in \mathfrak{C}$ such that $a \not\downarrow_M^s b$, then $\text{tp}(b/Ma)$ forks over M .*

Proof. Otherwise, we could find a global M -invariant extension \tilde{p} of $\text{tp}(b/Ma)$. Take $(a_i)_{i < |T|^+}$ to be a sequence of realizations of $\text{tp}(a/M)$ with $a_0 = a$ and $a_i \downarrow_M^s a_{<i}$ for each i . By invariance, if $b_* \models \tilde{p}$ over everything, for each $i < |T|^+$, $\text{tp}(b_*, a_i/M) = \text{tp}(b_*, a/M)$ and $b_* \not\downarrow_M^s a_i$. This contradicts Proposition 3.23. \square

Corollary 3.26. *Let a and b be distant over M , then $\text{tp}(a/Mb)$ forks over M if and only if $\text{tp}(b/Ma)$ forks over M if and only if $a \not\downarrow_M b$.*

Proposition 3.27. *Let $p \in S(M)$ be an invariant type and $q \in S(M)$ be distant from p . Let $I = (a_i)_{i < \omega}$ be a Morley sequence of p over M and $b \models q$. Then $\lim(I/Mb) = p|_{Mb}$.*

Proof. This follows easily from Proposition 3.23 by making the sequence I of large cardinality. \square

EXAMPLE 3.28 (ACVF). Take T to be ACVF, and M a model of T . Let $p \in S(M)$ be an invariant type of a field element. By [6], Corollary 12.14, there are definable functions f and g respectively into the residue field k and the value group Γ such that letting $p_k = f_*(p)$ and $p_\Gamma = g_*(p)$, we have:

For any $a \models p$ and $b \in \mathfrak{C}$, $\text{tp}(a/Mb) = p|_{Mb}$ if and only if $\text{tp}(f(a)/Mb) = p_k|_{Mb}$ and $\text{tp}(g(a)/Mb) = p_\Gamma|_{Mb}$.

Take such an invariant type p and $a \models p$. Then a is s -dominated by $f(a)$ since if $b \in \mathfrak{C}$ is distant from a over M , then by distality of Γ , $\text{tp}(b/M)$ and $\text{tp}(g(a)/M)$ are weakly orthogonal.

3.2 The finite-co-finite theorem and application

We prove now an analog of Proposition 3.23 which does not require to work over a model. We prove it by reproducing the proof of that proposition in the context of domination for indiscernible sequences.

Proposition 3.29. *Let A be any set of parameters and let p be some global A -invariant type. Let $a \in \mathfrak{C}$. Let I be an infinite Morley sequence of p over Aa and J be an infinite Morley sequence of p over AI . Let $\phi(x; y) \in L(A)$, then the set $\{b \in J : \models \phi(b, a)\}$ is finite or co-finite in J .*

Proof. Assume not. Then we may expand I to a saturated sequence. Without loss, the formula $\phi(x, b)$ is true for $x \in I$ and pruning J , we may assume that it is false for $x \in J$. Finally, we may expand J so that $J = \langle b_i : i < |T|^+ \rangle$.

We can find sequences $\langle \bar{b}_*^i : i < |T|^+ \rangle$ such that :

- Each \bar{b}_*^i fills some cut of I , the cuts being distinct from one another, and the \bar{b}_*^i are placed independently over I ;

- for each index i , \bar{b}_*^i strongly dominates b_i over (I, A) .

(Why ? First take \bar{d}_*^0 strongly dominating b_0 over (I, A) . Let $\langle b'_i : 0 < i < |T|^+ \rangle$ be a Morley sequence of p over everything. There is an automorphism σ fixing AIb_0 sending $\langle b'_i : 0 < i < |T|^+ \rangle$ to $\langle b_i : 0 < i < |T|^+ \rangle$. Let $\bar{b}_*^0 = \sigma(\bar{d}_*^0)$. Then take \bar{d}_*^1 strongly dominating b_1 over (I, A) with $\bar{d}_*^1 \downarrow_I \bar{b}_*^0$. And iterate.)

Let I' be the sequence I with all the \bar{b}_*^i added in their respective cuts. It is an A -indiscernible sequence. By shrinking of indiscernibles, there is $I'' \subseteq I$ obtained by removing at most $|T|$ of the tuples \bar{b}_*^i from I' such that I'' is indiscernible over Aa . Without loss, assume we have not removed the tuple \bar{b}_*^0 . Then by Proposition 3.7 (External characterization), $b_0 \models p|_{Aa}$. This contradicts the hypothesis. \square

Theorem 3.30 (Finite-co-finite theorem). *Let $I = I_1 + I_2 + I_3$ be indiscernible, I_1 and I_3 being infinite. Assume that $I_1 + I_3$ is A -indiscernible and take $\phi(x; a) \in L(A)$, then the set $B = \{b \in I_2 : \models \phi(b; a)\}$ is finite or co-finite.*

Proof. This follows from the previous proposition by setting p to be the limit type of I_3^* (I_3 in reverse order). \square

Note that necessarily, B in the statement of the theorem is finite if $\neg\phi(b; a)$ holds for $b \in I_1 + I_3$ and co-finite otherwise (because you can incorporate some parts of I_1 and I_3 to I_2 , also it follows from the proof). This will be used implicitly in applications.

Corollary 3.31. *Let $I = I_1 + I_2 + I_3$ be indiscernible, I_1 and I_3 being infinite with no endpoints and I_2 densely ordered. Assume that $I_1 + I_3$ is A -indiscernible. Write $I_2 = (a_i)_{i \in \mathcal{I}}$. Then given some linear order $\mathcal{J} \supseteq \mathcal{I}$, one can find tuples a_i , $i \in \mathcal{J} \setminus \mathcal{I}$ such that:*

- $I_1 + \langle a_i : i \in \mathcal{J} \setminus \mathcal{I} \rangle + I_3$ is indiscernible over A ,
- $I_1 + \langle a_i : i \in \mathcal{J} \rangle + I_3$ is indiscernible.

Proof. We construct the points a_i , $i \in \mathcal{J} \setminus \mathcal{I}$ simply by realizing limit types of cuts of I_2 over everything. More precisely, given \mathfrak{c} a cut of \mathcal{I} , identify \mathfrak{c} with the corresponding cut of I_2 . Assume for simplicity that \mathfrak{c} has infinite cofinality from the right and let $p_{\mathfrak{c}}$ be $\lim(\mathfrak{c}^+)$ (seen a global type). Note that if $\mathfrak{c} \neq \mathfrak{c}'$, then the types $p_{\mathfrak{c}}$ and $p_{\mathfrak{c}'}$ commute. Let $\mathcal{J}_{\mathfrak{c}}$ be the convex subset of \mathcal{J} formed by elements falling in the cut \mathfrak{c} . Finally take $\langle a_i : i \in \mathcal{J} \setminus \mathcal{I} \rangle$ to realize $\bigotimes_{\mathfrak{c}} p_{\mathfrak{c}}^{(\mathcal{J}_{\mathfrak{c}})}$ over IA .

The second condition is obviously satisfied, so we have to check the first one. We start by considering a cut \mathfrak{c} , and show that $I_1 + \langle a_i : i \in \mathcal{J}_{\mathfrak{c}} \rangle + I_3$ is indiscernible over A . The fact that for $i \in \mathcal{J}_{\mathfrak{c}}$, and $a \in I_1$, $\text{tp}(a_i/A) = \text{tp}(a/A)$ follows immediately from the finite-co-finite theorem 3.30. Now consider $i < j \in \mathcal{J}_{\mathfrak{c}}$ and $\phi(x_1, x_2) \in L(A)$ a formula. Assume that for $a \in I_1$, $b \in I_3$ we have $\models \phi(a, b)$.

Write $c = (\mathcal{K}_1, \mathcal{K}_2)$, seen as a cut of \mathcal{I} . By construction of $(a_i)_{i \in \mathcal{J}_c}$ and shrinking of indiscernibles (Proposition 1.10), we have:

$$\models \phi(a_i, a_j) \iff \text{for some coinitial } \mathcal{K} \subset \mathcal{K}_2, \phi(a_s, a_t) \text{ holds for } s < t \in \mathcal{K}.$$

Assume we have $\neg \phi(a_i, a_j)$. So easily, we can find points $s_1 < t_1 < s_2 < t_2 < \dots \in \mathcal{K}_2$ such that $\neg \phi(a_{s_k}, a_{t_k})$ holds for each $k < \omega$. Let $L_2 = \langle a_{s_k} \hat{~} a_{t_k} : k < \omega \rangle$. Take also L_1 to be any sequence of increasing pairs of members of I_1 , so that $L_1 + L_2$ is indiscernible, and pick similarly L_3 . Then the finite-co-finite theorem applied to the sequence $L_1 + L_2 + L_3$ gives us a contradiction.

We can do the same reasoning if $\phi(x_1, x_2)$ has parameters in AI_1I_2 (by adding parts of I_1I_2 to A and decreasing them). Also one sees at once that the construction generalizes to formulas $\phi(x_1, \dots, x_n)$ with more variables and we obtain that $I_1 + \langle a_i : i \in \mathcal{J}_c \rangle + I_3$ is indiscernible over A .

Next, we look at two cuts $\mathfrak{c}_1 < \mathfrak{c}_2$ and we want to see that $I_1 + \langle a_i : i \in \mathcal{J}_{\mathfrak{c}_1} + \mathcal{J}_{\mathfrak{c}_2} \rangle + I_3$ is indiscernible over A . We know that $\langle a_i : i \in \mathcal{J}_{\mathfrak{c}_2} \rangle$ realizes $p_{\mathfrak{c}_2}^{(\mathcal{J}_{\mathfrak{c}_2})}$ over everything else. We may assume that $\mathcal{J}_{\mathfrak{c}_1}$ is without endpoints. Take some finite $\mathcal{K}_0 \subset \mathcal{J}_{\mathfrak{c}_1}$ and let \mathcal{K}_1 be $\{i \in \mathcal{J}_{\mathfrak{c}_1} : i > \mathcal{K}_0\}$. Then the sequence $\langle a_i : i \in \mathcal{K}_1 \rangle + I_3$ is indiscernible over $A \cup \{a_i : i \in \mathcal{K}_0\}$. The same reasoning as above shows that the sequence $\langle a_i : i \in \mathcal{K}_1 \rangle + \langle a_i : i \in \mathcal{J}_{\mathfrak{c}_2} \rangle + I_3$ is indiscernible over $A \cup \{a_i : i \in \mathcal{K}_0\}$. It follows that $I_1 + \langle a_i : i \in \mathcal{J}_{\mathfrak{c}_1} + \mathcal{J}_{\mathfrak{c}_2} \rangle + I_3$ is indiscernible over A .

Iteratively, we prove that $I_1 + \langle a_i : i \in \mathcal{J}_{\mathfrak{c}_1} + \dots + \mathcal{J}_{\mathfrak{c}_n} \rangle + I_3$ is indiscernible over A and finally, that $I_1 + \langle a_i : i \in \mathcal{J} \setminus \mathcal{I} \rangle + I_3$ is indiscernible over A . \square

Corollary 3.32. *Let $I_1 + I_2 + I_3$ be an indiscernible sequence of finite tuples, with I_1 and I_3 infinite without endpoints. Assume that $I_1 + I_3$ is indiscernible over A . Then we can find some subsequence $I'_2 \subset I_2$ with $I_2 \setminus I'_2$ of size at most $|T| + |A|$ such that $I_1 + I'_2 + I_3$ is indiscernible over A .*

Proof. Without loss, we may assume that I_2 is densely ordered. Write $I_2 = \langle a_i : i \in \mathcal{I} \rangle$ and take some $|\mathcal{I}|^+$ -saturated linear order $\mathcal{J} \supset \mathcal{I}$. By Corollary 3.31 we can find tuples $\langle a_i : i \in \mathcal{J} \setminus \mathcal{I} \rangle$ such that :

- $I_1 + \langle a_i : i \in \mathcal{J} \setminus \mathcal{I} \rangle + I_3$ is indiscernible over A ,
- $I_1 + \langle a_i : i \in \mathcal{J} \rangle + I_3$ is indiscernible.

By shrinking of indiscernibles, there is $\mathcal{J}_0 \subset \mathcal{J}$ of size at most $|T| + |A|$ such that $I_1 + \langle a_i : i \in \mathcal{J} \setminus \mathcal{J}_0 \rangle + I_3$ is indiscernible. Then set $I'_2 = \langle a_i : i \in \mathcal{I} \setminus \mathcal{J}_0 \rangle$. \square

We now give an application of this result to externally definable sets.

We will use the following notation: if $M \models T$, $M < N$ is an elementary extension and $A \subseteq N$ containing M , then $M_{[A]}$ is the structure with universe

M with language composed of a predicate for every subset of M^l (any l) of the form $\phi(M; \bar{c})$, $\bar{c} \in A^k$ for any $\phi(\bar{x}; \bar{y}) \in L(M)$, interpreted the obvious way.

Shelah proved in [12] that $M_{[\mathfrak{c}]}$ eliminates quantifiers. We refer the reader to [3] for a slightly different approach, that we will use (and recall) here. If $p \in S(M)$ is any type and $a \models p$, then it is not true in general that $M_{[a]}$ eliminates quantifiers (see [3], Example 1.8 for a counterexample). However it is conjectured in [3] that $M_{[I]}$ does, where I is a coheir sequence starting with a . We prove a special case of this when p is interior to M . See the definition below.

We will need some notions from [3] that we recall now. If X is an externally definable subset of X (i.e., a subset of the form $\phi(M, c)$ for some tuple $c \in \mathfrak{C}$), then an *honest definition* of X is a formula $\theta(x, d) \in L(\mathfrak{C})$ such that (1) $\theta(M, d) = X$ and (2) for every formula $\psi(x) \in L(M)$ such that $X \subseteq \psi(M)$ then $\mathfrak{C} \models \theta(x) \rightarrow \psi(x)$.

Lemma 3.33. *If $A \subset \mathfrak{C}$ containing M is such that for every formula $\phi(x; c) \in L(A)$, $\phi(M; c)$ has an honest definition with parameters in A , then $M_{[A]}$ eliminates quantifiers.*

Proof. Let $\phi(x, y; c) \in L(A)$ and let $\theta(x, y; d) \in L(A)$ be an honest definition of $X := \phi(M; c)$. Let $\pi : M^{|x|+|y|} \rightarrow M$ be the projection on the first $|x|$ coordinates. Let $\psi(x; d) = \exists y \theta(x, y; d)$. Then $\psi(M; d) = \pi(X)$: it is clear that $\psi(M; d) \subseteq \pi(X)$, and if $a \in M^{|x|} \setminus \pi(X)$, then the set $\{(x, y) \in M^{|x|+|y|} : y \neq a\}$ contains X and by honesty $\mathfrak{C} \models \theta(x, y; d) \rightarrow y \neq a$ which gives the reverse inclusion. \square

Definition 3.34. Let p be an M -invariant global type. We say that p is *interior* to M if $p^{(\omega)}$ is both an heir and a co-heir of its restriction to M .

An example of an interior type is given by the following situation: let $I \subset M$ be indiscernible and \mathfrak{c} a cut interior to I such that M respects \mathfrak{c} . Then the type $p = \lim(\mathfrak{c}^+)$ is interior to M .

Lemma 3.35. *Let p be a global M -invariant type interior to M . Let $I_0 + I_1 + I_2$ be a Morley sequence of p over M . For $i < 3$ let $\bar{a}_i \subset I_i$ be a finite tuple. Assume that $\bar{a}_1 \models \phi(\bar{x}; \bar{a}_0, \bar{a}_2)$, $\phi \in L(M)$, then there are two tuples $\bar{b}_0, \bar{b}_2 \subset M$ such that $\bar{a}_1 \models \phi(\bar{x}; \bar{b}_0, \bar{b}_2)$.*

Proof. First find \bar{b}_2 such that $\bar{a}_1 \models \phi(\bar{x}; \bar{a}_0, \bar{b}_2)$ by the coheir hypothesis. Then find \bar{b}_0 by the heir hypothesis. \square

Theorem 3.36 (Shelah expansion for interior types). *Let p be a global M -invariant type interior to M . Let I be a Morley sequence of p over M . Then $M_{[I]}$ eliminates quantifiers.*

Proof. Take a saturated extension $M_{[I]} < N^*$ of size $\kappa > |M|$. The model N^* can be seen as a reduct to the language of $M_{[I]}$ of some $N_{[J]}$ for $M < N$ and $J \equiv_M I$, J indiscernible over N . Without loss $I = J$. Notice that N^* and $N_{[I]}$ have the same definable sets.

Claim: There is an indiscernible sequence $I_1 + I_2 \subset N$ such that N respects the cut $c = (I_1, I_2)$ and $I \models \lim(c^+)^{(\omega)}$.

Proof: Write $N = \bigcup_{i < \kappa} A_i$ with $|A_i| < \kappa$. Let $i < \kappa$. By Lemma 3.35 and saturation, we can find sequences $K_i, L_i \subset N$ of order type ω such that $K_i + I + L_i$ is indiscernible over A_i . Let $I_1 = K_1 + K_2 + \dots$ and $I_2 = \dots + L_2 + L_1$, the sums ranging over $i < \kappa$. The required property is then easy to check.

Let $\phi(x; y)$ be a formula and $a_0 \models p$, $a_0 \in I$. We consider the pair (M, N) and show that $\phi(a_0; M)$ has an honest definition with parameters in $M + I_1 + I_2$.

By the Theorem 3.30 and compactness, there are integers n, N and a finite set of formulas δ such that for every finite sequence $J_1 + J_3 + J_2$, satisfying:

- J_1 and J_2 are of size at least n ,
- $J_1 + J_3 + J_2$ is indiscernible,
- $J_1 + J_2$ is δ -indiscernible over b and
- $\phi(x; b)$ holds on all elements of J_1 and J_2 ,

then $\neg\phi(x; b)$ holds on at most N elements of J_3 .

Let $I'_1 \subset I_1$ and $I'_2 \subset I_2$ be finite of size n such that $I'_1 + I'_2$ is M -indiscernible. Consider the formula $\theta(y) \in L(MI)$ such that if $b \models \theta(y)$, then $I'_1 + I'_2$ is δ -indiscernible over a , and $\phi(\bar{a}_0; y)$ holds on all elements of $I'_1 + I'_2$. Define analogously $\theta_1(y)$ using $\neg\phi$ instead of ϕ .

Then, for every $b \in M$, $\theta(b)$ holds if and only if $\phi(a_0; b)$ holds. Also, if $b \in N$, and $\theta(b)$ holds, then $\phi(a_0; b)$ holds (Why? Only finitely many elements a from $I_1 + I_2$, with $I'_1 < a < I'_2$ can satisfy $\phi(a; b)$). This easily implies that θ is an honest definition of $\phi(a_0; M)$.

To conclude the theorem, notice that we can do the same thing replacing p by $p^{(n)}$ for any n , which takes care of formulas $\phi(\bar{a}; y)$ with \bar{a} a finite subset of I instead of one element. \square

4 Sharp theories

In this last section, we study theories in which types are s-dominated by generically stable types. We show that this is implied by the existence of some form of decomposition of indiscernible sequences into “stable by distal”. Our goal is to give a criterion which we can check in dimension 1 and conclude that dp-minimal theories are sharp. One could probably introduce stronger notions, and ask for example that types are s-dominated by types living in a stable sort, but we do not pursue this here.

Definition 4.1. The theory T is sharp if for every $|T|^+$ -saturated model M and $p \in S(M)$ an invariant type realized by a , there is some generically stable type $q \in S(M)$ and $a_* \models q$ such that a_* s-dominates a over M .

Definition 4.2. Let $I = \langle a_i : i \in \mathcal{I} \rangle$ be a dense indiscernible sequence. A decomposition of I is an indiscernible sequence $K = \langle a_i \hat{\ } b_i : i \in \mathcal{I} \rangle$ where the sequence $J = (b_i)_{i \in \mathcal{I}}$ is totally indiscernible and such that:

For every K' of same EM-type as K , \mathfrak{c} a Dedekind cut of K' , $d \in \mathfrak{C}$ such that K' is indiscernible over d and $a \hat{\ } b$ filling \mathfrak{c} ; if there is a' such that $a' \hat{\ } b$ fills \mathfrak{c} over dK' , then $a \hat{\ } b$ fills \mathfrak{c} over dK' .

By usual sliding argument, if K is dense and contains some Dedekind cut \mathfrak{c} , it is enough to check the condition for $K' = K$.

An indiscernible sequence $I = \langle a_i : i \in \mathcal{I} \rangle$ is *decomposable* if it admits a decomposition $K = \langle a_i \hat{\ } b_i : i \in \mathcal{I} \rangle$. In this case, we will say that I is decomposable over $\langle b_i : i \in \mathcal{I} \rangle$.

Remark 4.3. There are two trivial cases of decomposability: If I is distal, then it is decomposable over the sequence of empty tuples, if I is totally indiscernible, it is decomposable over itself.

Lemma 4.4 (Internal characterization). *An indiscernible sequence of pairs $I = (a_i \hat{\ } b_i)_{i \in \mathcal{I}}$ is a decomposition, if and only if the following holds:*

☒ *For every J, K, L infinite indiscernible sequences without endpoints of same EM-type as I and $a \hat{\ } b$, $a' \hat{\ } b'$, if $J + a \hat{\ } b + K + L$, $J + K + a' \hat{\ } b' + L$ are indiscernible, and there exist a_0, a'_0 such that $J + a_0 \hat{\ } b + K + a'_0 \hat{\ } b' + L$ is indiscernible, then $J + a \hat{\ } b + K + a' \hat{\ } b' + L$ is indiscernible.*

Proof. Assume that I is a decomposition. Then taking $d = a'_0 \hat{\ } b' + L$ in the definition, we see that $J + a \hat{\ } b + K$ is indiscernible over $a'_0 \hat{\ } b' + L$. Then taking $d =$

$J+a\hat{b}$, we get that $K+a'\hat{b}'+L$ is indiscernible over $J+a\hat{b}$, so $J+a\hat{b}+K+a'\hat{b}'+L$ is indiscernible.

Conversely, assume \boxtimes holds and without loss \mathcal{I} is a dense order. Notice that the analog of \boxtimes where we fill n cuts instead of 2 follows from \boxtimes by easy induction (as in Lemma 2.6). Let $d \in \mathfrak{C}$, \mathfrak{c} , $a\hat{b}$ and a' be as in the definition of decomposition. Assume that $a\hat{b}$ does not fill \mathfrak{c} over Ad . Adding parameters to d if necessary, we may assume that for some formula $\phi(x, y)$, and all $a_*\hat{b}_* \in I$, we have $\phi(a_*\hat{b}_*, d) \wedge \neg\phi(a\hat{b}, d)$. Fix some increasing sequence $(\mathfrak{c}_k)_{k < \omega}$ of Dedekind cuts of I . For each $k < \omega$, we can find a_k, a'_k, b_k such that $\text{tp}(a_k, a'_k, b_k, d/I)$ is similar to $\text{tp}(a, a', b, d/I)$ and $a_k\hat{b}_k$ fills the cut \mathfrak{c}_k . By \boxtimes and the remark above, the sequence obtained by adding all the tuples $a_k\hat{b}_k$ to I in their respective cuts is indiscernible. Then the formula $\phi(x, y)$ has infinite alternation rank. \square

We will need the following strengthening of Lemma 2.8.

Lemma 4.5 (Strong base change 2). *Let $I = (a_i\hat{b}_i)_{i \in \mathcal{I}}$ be an indiscernible sequence and $A \supset I$ a set of parameters. Let $(\mathfrak{c}_i)_{i \in \mathcal{J}}$ be a sequence of pairwise distinct polarized Dedekind cuts in I . Call \mathfrak{c}'_i the corresponding cut in the sequence $(b_i)_{i \in \mathcal{I}}$. For each i let $d_i\hat{e}_i$ fill the cut \mathfrak{c}_i . Assume also that the sequence $(e_i)_{i \in \mathcal{J}}$ realizes $\otimes \lim(\mathfrak{c}'_i)$ over I . Then there exist $(d'_i\hat{e}'_i)_{i \in \mathcal{J}}$ such that*

- $\text{tp}(\langle d'_i\hat{e}'_i : i \in \mathcal{J} \rangle / I) = \text{tp}(\langle d_i\hat{e}_i : i \in \mathcal{J} \rangle / I)$,
- for each i , $\text{tp}(d'_i\hat{e}'_i / A) = \lim(\mathfrak{c}_i / A)$,
- $(e'_i)_{i \in \mathcal{J}}$ realizes $\otimes_i \lim(\mathfrak{c}'_i)$ over A .

Proof. The proof is essentially the same as that of Lemma 2.8.

Assume the result does not hold. Then by compactness, we may assume that $\mathcal{J} = \{1, \dots, n\}$ and that there is a formula $\phi(x_1\hat{y}_1, \dots, x_n\hat{y}_n) \in \text{tp}(\langle d_i\hat{e}_i : i \rangle / I)$, a formula $\theta(y_1, \dots, y_n) \in \otimes \lim(t'_i/m)$ and formulas $\psi_i(x_i, y_i) \in \lim(\mathfrak{c}_i/m)$ for some finite $m \in A$ such that $\phi(x_1\hat{y}_1, \dots, x_n\hat{y}_n) \wedge \theta(y_1, \dots, y_n) \wedge_i \psi_i(x_i, y_i)$ is inconsistent. Let I_0 denote the parameters of ϕ .

Assume for simplicity that $n = 2$ (the proof for $n > 2$ is the same) and without loss \mathfrak{c}_i is polarized as \mathfrak{c}_i^- . For $i = 1, 2$ take $(I_i, I'_i) \trianglelefteq \mathfrak{c}_i$ such that ψ_i holds on all elements of I_i , $\theta(y_1, y_2)$ holds for each $(x_1\hat{y}_1, x_2\hat{y}_2) \in I_1 \times I_2$, and $I_i \cup I'_i$ contains no element of I_0 . Then $I_1 + I'_1$ and $I_2 + I'_2$ are mutually indiscernible over I_0 . So for every two cuts \mathfrak{d}_1 and \mathfrak{d}_2 respectively from $I_1 + I'_1$ and $I_2 + I'_2$, we can find points $d_1\hat{e}_1$ and $d_2\hat{e}_2$ filling those cuts (even seen as cuts of I) such that $\phi(d_1, e_1, d_2, e_2)$ holds and there are d'_1, d'_2 such that $(d'_1\hat{e}_1, d'_2\hat{e}_2)$ fills the polycut $(\mathfrak{d}_1, \mathfrak{d}_2)$ over I .

Take a cut \mathfrak{d}_1 inside I_1 and \mathfrak{d}_2 inside I_2 and see them as cuts of I . Fill \mathfrak{d}_1 by $d_1 \hat{\ } e_1$ and \mathfrak{d}_2 by $d_2 \hat{\ } e_2$ as above. By hypothesis, either $\neg\theta(e_1, e_2)$, $\neg\psi_1(d_1, e_1)$ or $\neg\psi_2(d_2, e_2)$ holds. In one of the latter two cases, proceed as in Lemma 2.8. In the first case, keep e_1 and e_2 and add points (d'_1, d'_2) such that I with $d'_1 \hat{\ } e_1$ and $d'_2 \hat{\ } e_2$ added is indiscernible. Then iterate with $I \cup \{d'_1 \hat{\ } e_1, d'_2 \hat{\ } e_2\}$ instead of I .

After iterating this ω times, either ψ_1 , ψ_2 or θ has infinite alternation rank. \square

Lemma 4.6 (Base change). *The notion of being a decomposition is stable both ways under base change: If $(a_i \hat{\ } b_i)_{i \in \mathcal{I}}$ is A -indiscernible, then it is a decomposition in T if and only if it is a decomposition in $T(A)$.*

Proof. Assume $I = (a_i \hat{\ } b_i)_{i \in \mathcal{I}}$ is a decomposition, then it follows immediately from the definition that it is a decomposition from the point of view of $T(A)$.

For the converse, use the internal characterization and strong base change 2 (Lemma 4.5) as in the proof of Corollary 2.9. \square

Lemma 4.7. *Assume that $K = (a_i \hat{\ } b_i)_{i \in \mathcal{I}}$ is a decomposition of $I = (a_i)_{i \in \mathcal{I}}$. Let \mathfrak{c} be a cut of K filled by a sequence L and denote by \mathfrak{c}' the corresponding cut in $(b_i)_{i \in \mathcal{I}}$. Let L_2 the projection on L on the second factor (so L_2 is a totally indiscernible sequence). Let $d \in \mathfrak{C}$ be such that K is indiscernible over d and L_2 is a Morley sequence of the limit type of \mathfrak{c}' over Kd . Then $K \cup L$ is indiscernible over d (where L is placed in the cut \mathfrak{c}).*

Proof. Assume L is dense of size $|T|$ and using Corollary 3.31, increase L to some saturated sequence L' filling \mathfrak{c} and such that the sequence $K_0 = K \cup (L' \setminus L)$ is indiscernible over d . Let now $a_1 \hat{\ } b_1 \in L$. This tuple fills a Dedekind cut of K_0 . By domination in the sequence K_0 , we see that $K_1 = K_0 \cup \{a_1 \hat{\ } b_1\}$ is indiscernible over d . Then we can take some other $a_2 \hat{\ } b_2 \in L$. It fills a Dedekind cut of K_1 and by domination in K_1 , $K_2 = K_1 \cup \{a_2 \hat{\ } b_2\}$ is indiscernible over d . Iterating, we see that $K \cup L'$ is indiscernible over d and therefore $K \cup L$ is indiscernible over d . \square

Lemma 4.8. *Let M be a $|T|^+$ -saturated model and $p, q \in S(M)$ be two commuting invariant types. Take $I \models p^{(\omega)}$ and any $b \models q$. Then we may find two sequences I_1, I_2 such that $I_1 + I + I_2$ is a Morley sequence of p over M and $I_1 + I_2$ is a Morley sequence of p over Mb .*

Proof. Let r be the inverse of p over M (recall the definition as stated after Lemma 2.19). We take I_2 to be a Morley sequence of p over MIb and then I_1 to be a Morley sequence of r , indexed in the opposite order, over MII_2b . Over M , the Morley sequence of r is the opposite of the Morley sequence of p so the first statement follows. To see the second statement, recall that if $s \in S(M)$ is any invariant type, then $r_x \otimes s_y|_M = s_y \otimes p_x|_M$. In particular,

$$r_x \otimes (q_y \otimes p_{x_1, \dots, x_n}^{(n)})|_M = (q_y \otimes p_{x_1, \dots, x_n}^{(n)}) \otimes p_x|_M = q_y \otimes p_{x, x_1, \dots, x_n}^{(n+1)}|_M.$$

The result follows. □

Proposition 4.9. *Assume that all sequences are decomposable, then T is sharp.*

Proof. Let M be $|T|^+$ -saturated and $p \in S(M)$ be an A -invariant type. Let $a \models p$. Let $I \subset M$ be a small dense Morley sequence of p over A and let $K \subset M$ be a decomposition of I . Let \mathfrak{c} be a Dedekind cut of K and \mathfrak{c}_1 the corresponding cut of I . As in the proof of the moving away lemma 3.15, construct some dense sequence \bar{d} realizing a power of $\lim(\mathfrak{c}_1^+/M)$ and such that \bar{d} s -dominates a over M . Extend \bar{d} to \bar{c} realizing a power of $\lim(\mathfrak{c}^+/M)$. So \bar{c} is the union of \bar{d} and some totally indiscernible sequence \bar{e} . The type of \bar{e} over M is generically stable.

Claim : \bar{e} s -dominates a over M .

Proof : Let $u \in \mathfrak{C}$ be distant from a and independent from \bar{e} over M . Let u_* realize an invariant type distant from $a\bar{c}$ over M such that u_* s -dominates u and is independent from \bar{c} over M . If we show that $u_* \downarrow_M \bar{d}$, then as \bar{d} s -dominates a it will follow that $u_* \downarrow_M a$ and therefore $u \downarrow_M a$. Replacing u by u_* , we may now assume that u is distant from $a\bar{c}$ over M and realizes an invariant type.

Call $r = \lim(\mathfrak{c}^+)$ (a global invariant type). By Lemma 4.8, let I_1 and I_2 be two sequences such that $I_1 + I_2$ is indiscernible over Mu and $I_1 + \bar{c} + I_2$ is indiscernible over M . Also as u is independent from \bar{e} over M , the hypothesis of Lemma 4.7 are satisfied (where L_2 there is \bar{e} here). We conclude that u is independent from \bar{d} over M and therefore u is independent from a over M . □

4.1 Reduction to dimension 1

We prove here the following proposition.

Proposition 4.10. *Assume that all sequences of elements are decomposable, then every sequence is decomposable.*

Assume from now on that all indiscernible sequences of elements of \mathfrak{C} are decomposable. We will take an arbitrary indiscernible sequence and build a decomposition for it adjoining totally indiscernible sequences to it one-by-one. The proof is an adaptation of the one from Section 2.4. We start with a base set of parameters A that we allow to grow freely during the construction. In what follows, we work over A , even when not explicitly mentioned. We have an indiscernible sequence $I = \langle a_i \hat{\alpha}_i : i \in \mathcal{I} \rangle$, where $\mathcal{I} = (0, 1)$ for simplicity and such that the sequence $\langle \alpha_i : i \in \mathcal{I} \rangle$ is totally indiscernible.

For every $i \in \mathcal{I}$, call \mathfrak{c}_i the cut “ i^+ ” of I and \mathfrak{c}'_i the associated cut in the sequence $\langle \alpha_i : i \in \mathcal{I} \rangle$.

Step 1 : Derived sequence

Assume we have a witness of non-decomposition in the following form :

- A tuple $b \in \mathfrak{C}$, some $j \in (0, 1)$ and a pair (a, α) such that :
- $a \hat{\alpha}$ fills the cut \mathfrak{c}_j of I ,
- I is b -indiscernible,
- α realizes the type $\lim(\mathfrak{c}'_j)$ over Ib ,
- I with $a_j \hat{\alpha}_j$ replaced by $a \hat{\alpha}$ is not indiscernible over b .

As in Section 2.4, adding parameters to the base, we may assume that b is a single point, and that $\text{tp}(a \hat{\alpha}/b) \neq \text{tp}(a_j \hat{\alpha}_j/b)$. Let $r = \text{tp}(a \hat{\alpha}, b)$.

We construct a new sequence $\langle a'_i \hat{\alpha}'_i : i \in \mathcal{I} \rangle$ such that :

- $a'_i \hat{\alpha}'_i$ fills the cut \mathfrak{c}_i of I ,
- $\text{tp}(a'_i \hat{\alpha}'_i, b) = r$ for each i ,
- The sequence $(\alpha'_i)_{i \in \mathcal{I}}$ realizes $\bigotimes_{i \in \mathcal{I}} \lim(\mathfrak{c}'_i)$ over Ib ,
- The sequence $\langle a_i \hat{\alpha}_i \hat{a}'_i \hat{\alpha}'_i : i \in \mathcal{I} \rangle$ is b -indiscernible.

This is possible by indiscernibility of $(a_i \hat{\alpha}_i)_i$ over b (first pick the points α'_i then choose the a_i filling the cuts and then extract).

Step 2 : Constructing an array

Using Lemma 4.5, iterate this construction to obtain an array $\langle a_i^n \hat{\alpha}_i^n : i \in \mathcal{I}, n < \omega \rangle$ and sequence $\langle b_n : n < \omega \rangle$ such that :

- $a_i^0 \hat{\alpha}_i^0 = a_i \hat{\alpha}_i$ for each i ,
- For each $i \in \mathcal{I}$, $0 < n < \omega$, the tuple $a_i^n \hat{\alpha}_i^n$ realizes $\lim(\mathbf{c}_i)$ over $\langle b_k, a_i^k \hat{\alpha}_i^k : i \in \mathcal{I}, k < n \rangle$,
- For each $0 < n < \omega$, the sequence $(\alpha_i^n)_{i \in \mathcal{I}}$ realizes the type $\bigotimes_{i \in \mathcal{I}} \lim(\mathbf{c}'_i)$ over $\langle b_k, a_i^k \hat{\alpha}_i^k : i \in \mathcal{I}, k < n \rangle$,
- For each $0 < n < \omega$, $tp(b_n, \langle a_i^n \hat{\alpha}_i^n : i \in \mathcal{I} \rangle / I) = tp(b, \langle a'_i \hat{\alpha}'_i : i \in \mathcal{I} \rangle / I)$.

Claim : For every $\eta : \mathcal{I}_0 \subset \mathcal{I} \rightarrow \omega$ injective, the sequence $\langle a_i^{\eta(i)} \hat{\alpha}_i^{\eta(i)} : i \in \mathcal{I}_0 \rangle$ is indiscernible, of same EM-type as I .

The sequence $U = \langle \alpha_i^n : (i, n) \in \mathcal{I} \times \omega \rangle$, where $\mathcal{I} \times \omega$ is ordered lexicographically, is totally indiscernible.

Proof. Easy, by construction. \square

Expanding and extracting, we may assume that the sequence of rows $\langle b_n + (a_i^n \hat{\alpha}_i^n)_{i \in \mathcal{I}} : 0 < n < \omega \rangle$ is indiscernible. By assumption all sequences of points are decomposable. So let $(b_n \hat{\beta}_n)_{n < \omega}$ be an decomposition of $(b_n)_{n < \omega}$. Expanding and extracting again, we may assume that the new sequence of rows $\langle b_n \hat{\beta}_n + (a_i^n \hat{\alpha}_i^n)_{i \in \mathcal{I}} : 0 < n < \omega \rangle$ is indiscernible and that the sequence of columns $\langle (a_i^n \hat{\alpha}_i^n)_{0 < n < \omega} : i \in \mathcal{I} \rangle$ is indiscernible over $\{b_n \hat{\beta}_n : n < \omega\}$.

Step 3 : Conclusion

Claim : The sequences $(b_n \hat{\beta}_n)_{n < \omega}$ and $\langle (a_i^n \hat{\alpha}_i^n)_{i \in \mathcal{I}} : 0 < n < \omega \rangle$ are weakly linked.

The sequences $(b_n \hat{\beta}_n)_{n < \omega}$ and U are mutually indiscernible.

Proof. For the first statement, the proof is the same is in Section 2.4.

The second statement is similar. If for example we have $\phi(b_n, \beta_n, \alpha_i^n)$, then $\phi(b_n, \beta_n, \alpha_j^n)$ must hold for all $j \in \mathcal{I}$, and therefore by total indiscernibility of U and NIP , $\phi(b_n, \beta_n, \alpha_j^m)$ must hold for every $(j, m) \in \mathcal{I} \times \omega$. \square

Let $(c_n, \gamma_n) = (a_{1-\frac{1}{n}}^n, \alpha_{1-\frac{1}{n}}^n)$, then :

1. The sequence $(c_n \hat{\gamma}_n)_{n < \omega}$ is indiscernible, with same EM-type as I ;
2. The sequences $(\gamma_n)_{n < \omega}$ and $(b_n \hat{\beta}_n)_{n < \omega}$ are mutually indiscernible;

3. The sequences $(c_n \hat{\gamma}_n)_{n < \omega}$ and $(b_n \hat{\beta}_n)_{n < \omega}$ are weakly linked;
4. We have $\text{tp}(c_n \hat{\delta}_n, b_m) = r$ if and only if $n = m$.

Consider the indiscernible sequence $(c_n \hat{\gamma}_n \hat{b}_n \hat{\beta}_n)_{n < \omega}$. We may increase it to an indiscernible sequence $(c_n \hat{\gamma}_n \hat{b}_n \hat{\beta}_n)_{n \in \mathcal{I}}$. Take some $n_0 \in \mathcal{I}$ and set $\mathcal{I} = \mathcal{I}_1 + \{n_0\} + \mathcal{I}_2$. Then by point 3 above, the sequence $\langle b_n \hat{\beta}_n : n \in \mathcal{I}_1 + \mathcal{I}_2 \rangle$ is indiscernible over $c_{n_0} \hat{\gamma}_{n_0}$. Therefore point 4 and the definition of decomposition imply that β_{n_0} does not realize the limit type of $\langle \beta_n : n \in \mathcal{I}_1 \rangle$ over $\{b_n \hat{\beta}_n : n \in \mathcal{I}_1 + \mathcal{I}_2\} \cup \{c_{n_0} \hat{\gamma}_{n_0}\}$. Adding parameters to the base, we may assume that it does not realize that limit type over $c_{n_0} \hat{\gamma}_{n_0}$.

We then iterate the construction, starting with the sequence $(c_n \hat{\gamma}_n \hat{\beta}_n)_{n \in \mathcal{I}}$. Assume that we can do this $|T|^+$ steps. We have at the end some base set of parameters A , an A -indiscernible sequence $\langle c_n \hat{(\alpha_n^i : i < |T|^+)} : n < \omega \rangle$ (we replaced the index set \mathcal{I} by ω for convenience) such that for each $i < |T|^+$, the sequence $(\alpha_n^i)_{n < \omega}$ is totally indiscernible over $A \cup \{\alpha_n^j : n < \omega, j \neq i\}$ but not indiscernible over $A \cup \{c_n, n < \omega\} \cup \{\alpha_n^j : n < \omega, j < i\}$. By Fodor's lemma, removing some sequences $(\alpha_n^i)_{n < \omega}$ and adding them to A , we may assume that for every i , $(\alpha_n^i)_{n < \omega}$ is not indiscernible over $A \cup \{d_n, n < \omega\}$. But this contradicts Proposition 1.18.

Therefore this construction must stop after less than $|T|^+$ stages. At the end, we obtain a decomposition of the sequence we started with. This proves Proposition 4.10.

Corollary 4.11. *Every dp-minimal theory is sharp.*

EXAMPLE 4.12 (Non-sharp theory). Let L_0 be the language $\{R_n(x, y) : n < \omega\}$ and construct an L_0 structure M_0 as follows: the universe of M_0 is \mathbb{Q} , the ordinary rational numbers, and for every $x, y \in M_0$, $M_0 \models R_n(x, y)$ if and only if $x < y \wedge |x - y| < n$ holds in \mathbb{Q} . Let $T_0 = \text{Th}(M_0)$. Non-realized 1-types over M_0 satisfying $R_n(x, a)$ for some $n < \omega$ and $a \in M$ are in natural bijection with cuts of $(\mathbb{Q}, <)$. In addition to these, there is just one non-realized type $p \in S_1(M_0)$ which satisfies $\neg R_n(x, a)$ for every $n < \omega$ and $a \in M$. This type p is generically stable (and \emptyset -invariant). One can check easily that T_0 is dp-minimal.

Now consider $L_1 = L_0 \cup \{<\}$ where $<$ is a new binary relation. We expand M_0 to an L_1 -structure M_1 by making $<$ into a generic order (i.e., every L_0 -infinite definable set of M_1 is dense co-dense with respect to $<$. See for example [15]). A 1-type over M_1 is determined by its reduct to L_0 plus its $<$ -cut. Let $T_1 = \text{Th}(M_1)$. Easily, T_1 eliminates imaginaries so there are no generically

stable types (because the structure is linearly ordered). However T_1 is not distal: consider $I = (a_i)_{i \in \mathcal{I}}$ to be a dense $<$ -increasing sequence of points such that $\neg R_n(a_i, a_j)$ holds for every $n < \omega$ and $i, j \in \mathcal{I}$. Then this sequence is indiscernible and not distal. To see this, take two cuts \mathfrak{c}_1 and \mathfrak{c}_2 of I . Then there is a filling \mathfrak{c}_1 and b filling \mathfrak{c}_2 such that $R_1(a, b)$ holds. The generically stable type p in the reduct is detected by the non-distality of I .

We see however, that there is a natural ultra-imaginary stable sort: the quotient of M by the \forall -definable relation $E = \bigvee_{n < \omega} R_n$. And every point is in some sense s -dominated by its definable closure in that sort. It would be interesting to know if something like this is always true.

References

- [1] Hans Adler. An introduction to theories without the independence property. *Archive of Mathematical Logic*, 2008. to appear.
- [2] John Baldwin and Michael Benedikt. Stability theory, permutations of indiscernibles, and embedded finite models. *Transactions of the American Mathematical Society*, 352(11):4937–4969, 11 2000.
- [3] Artem Chernikov and Pierre Simon. Externally definable sets and dependent pairs. *Israel J. Math.*, 2012. DOI 10.1007/s11856-012-0061-9.
- [4] Artem Chernikov and Pierre Simon. Externally definable sets and dependent pairs ii. Submitted, 2012.
- [5] Alfred Dolich, David Lippel, and John Goodrick. Dp-minimal theories: basic facts and examples. *Notre Dame Journal of Formal Logic*, 52(3):267–288, 2011.
- [6] Deirdre Haskell, Ehud Hrushovski, and Dugald Macpherson. *Stable domination and independence in algebraically closed valued fields*, volume 30 of *Lecture Notes in Logic*. Association for Symbolic Logic, Chicago, IL, 2008.
- [7] Ehud Hrushovsk, Anand Pillay, and Pierre Simon. Generically stable and smooth measures in NIP theories. *Trans. Amer. Math. Soc.*, 2012. to appear.

- [8] Ehud Hrushovski, Ya'acov Peterzil, and Anand Pillay. Groups, measures, and the NIP. *J. Amer. Math. Soc.*, 21(2):563–596, 2008.
- [9] Ehud Hrushovski and Anand Pillay. On nip and invariant measures. *Journal of the European Mathematical Society*, 13:1005–1061, 2011.
- [10] Saharon Shelah. Strongly dependent theories. Submitted.
- [11] Saharon Shelah. Classification theory for elementary classes with the dependence property—a modest beginning. *Sci. Math. Jpn.*, 59(2):265–316, 2004. Special issue on set theory and algebraic model theory.
- [12] Saharon Shelah. Dependent first order theories, continued. *Israel Journal of Mathematics*, 173(1):1–60, 2009.
- [13] Saharon Shelah. A dependent dream and recounting types. *preprint*, 2010. 950 in Shelah archive.
- [14] Saharon Shelah. Dependent theories and the generic pair conjecture. *Communications in Contemporary Mathematics*, submitted, 2010. math.LO/0702292.
- [15] Saharon Shelah and Pierre Simon. Adding linear orders. *Journal of Symbolic Logic*, 77(2):717–725, 2012.
- [16] Pierre Simon. On dp-minimal ordered structures. *Journal of Symbolic Logic*, 76(4):448–460, 2011.
- [17] Pierre Simon. Finding generically stable measures. *Journal of Symbolic Logic*, 77(1):263–278, 2012.